arXiv:1403.2943v2 [math.NA] 21 Nov 2014

MULTILEVEL HYBRID CHERNOFF TAU-LEAP

ALVARO MORAES*, RAUL TEMPONE†, AND PEDRO VILANOVA‡

Abstract. In this work, we extend the hybrid Chernoff tau-leap method to the multilevel Monte Carlo (MLMC) setting. Inspired by the work of Anderson and Higham on the tau-leap MLMC method with uniform time steps, we develop a novel algorithm that is able to couple two hybrid Chernoff tau-leap paths at different levels. Using dual-weighted residual expansion techniques, we also develop a new way to estimate the variance of the difference of two consecutive levels and the bias. This is crucial because the computational work required to stabilize the coefficient of variation of the sample estimators of both quantities is often unaffordable for the deepest levels of the MLMC hierarchy. Our method bounds the global computational error to be below a prescribed tolerance, TOL, within a given confidence level. This is achieved with nearly optimal computational work. Indeed, the computational complexity of our method is of order $\mathcal{O}\left(TOL^{-2}\right)$, the same as with an exact method, but with a smaller constant. Our numerical examples show substantial gains with respect to the previous single-level approach and the Stochastic Simulation Algorithm.

Key words. Stochastic reaction networks, Continuous time Markov chains, Multilevel Monte Carlo, Hybrid simulation methods, Chernoff tau-leap, Dual-weighted estimation, Strong error estimation, Global error control, Computational Complexity

AMS subject classifications. 60J75, 60J27, 65G20, 92C40

1. Introduction.

This work, inspired by the multilevel discretization schemes introduced in [1], extends the hybrid Chernoff tau-leap method [19] to the multilevel Monte Carlo setting [9]. Consider a non-homogeneous Poisson process, X, taking values in the lattice of non-negative integers, $\mathbb{Z}_+^d$. We want to estimate the expected value of a given observable, $g : \mathbb{R}^d \to \mathbb{R}$ of X, at a final time, T, *i.e.*, $\mathrm{E}\left[g(X(T))\right]$. For example, in a chemical reaction in thermal equilibrium, the i-th component of X, $X_i(t)$, could describe the number of particles of species i present at time t. In the systems modeled here, different species undergo reactions at random times by changing the number of particles in at least one of the species. The probability of a single reaction happening in a small time interval is modeled by a non-negative propensity function that depends on the current state of the system. We present a formal description of the problem in Section 1.1.

Pathwise realizations of such pure jump processes (see, *e.g.*, [7]) can be simulated exactly using the Stochastic Simulation Algorithm (SSA), introduced by Gillespie in [10], or the Modified Next Reaction Method (MNRM) introduced by Anderson in [3]. Although these algorithms generate exact realizations for the Markov process, X, they are computationally feasible for only relatively low propensities.

For that reason, Gillespie in [11] and Aparicio and Solari in [4] independently proposed the tau-leap method to approximate the SSA by evolving the process with fixed time steps and by keeping the propensity fixed within each time step. In fact, the tau-leap method can be seen as a forward Euler method for a stochastic differential equation driven by Poisson random measures (see, *e.g.*, [17]).

*Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia (alvaro.moraesgutierrez@kaust.edu.sa).

†Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia (raul.tempone@kaust.edu.sa).

‡Computer, Electrical and Mathematical Sciences and Engineering Division, King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia (pedro.guerra@kaust.edu.sa).

A drawback of the tau-leap method is that the simulated process may take negative values, which is an undesirable consequence of the approximation and not a qualitative feature of the original process. For this purpose, we proposed in [19] a Chernoff-type bound that controls the probability of reaching negative values by adjusting the time steps. Also, to avoid extremely small time steps, we proposed switching adaptively between the tau-leap and an exact method, creating a hybrid tau-leap/exact method that combines the strengths of both methods.

More specifically, let $\bar{x}$ be the state of the approximate process at time t, and let $\delta \in (0,1)$ be given. The main idea is to compute a time step, $\tau{=}\tau(\delta,\bar{x})$, such that the probability that the approximate process reaches an unphysical negative value in $[t,t{+}\tau)$ is less than δ. This allows us to control the probability that a entire hybrid path exits the lattice, $\mathbb{Z}_+^d$. In turn, this quantity leads to the definition of the global exit error, which is a global error component along with the time discretization error and the statistical error (see Section 3.2 for details).

The multilevel Monte Carlo idea goes back at least to [13, 12]. In that setting, the main goal was to solve high-dimensional, parameter-dependent integral equations and to conduct corresponding complexity analyses. Later, in [9], Giles developed and analyzed multilevel techniques that were used to reduce the computational work when estimating an expected value using Monte Carlo path simulations of a certain quantity of interest of a stochastic differential equation. Independently, in [21], Speight introduced a multilevel approach to control variates. Control variates are a widespread variance reduction technique with the main goal of increasing the precision of an estimator or reducing the computational effort. The main idea is as follows: to reduce the variance of the standard Monte Carlo estimator of $\mathrm{E}[X]$,

$$\hat{\mu}_1 := \frac{1}{M}\sum_{m=1}^{M} X(\omega_m),$$

we consider another unbiased estimator of $\mathrm{E}[X]$,

$$\hat{\mu}_2 := \frac{1}{M}\sum_{m=1}^{M} \left(X(\omega_m) - (Y(\omega_m) - \mathrm{E}[Y])\right),$$

where Y is a random variable correlated with X with known mean, $\mathrm{E}[Y]$. The variable Y is called a control variate. Since $\mathrm{Var}[\hat{\mu}_2] {=} \mathrm{Var}[\hat{\mu}_1] {+} \mathrm{Var}[Y] {-} 2\mathrm{Cov}[X,Y]$, whenever $\mathrm{Cov}[X,Y] {>} \mathrm{Var}[Y]/2$, we have that $\mathrm{Var}[\hat{\mu}_2] {\leq} \mathrm{Var}[\hat{\mu}_1]$. If we assume that the computational work of generating the pair $(X(\omega),Y(\omega))$ is less than twice the computational work of generating $X(\omega)$, it is straightforward to conclude that $\hat{\mu}_2$ is preferred when $\rho^2_{X,Y} {>} 1/2$, where $\rho_{X,Y}$ is the correlation coefficient of the pair (X,Y). We observe that $\hat{\mu}_2$ can be written as

$$\hat{\mu}_2 = \mathrm{E}[Y] + \frac{1}{M}\sum_{m=1}^{M} (X-Y)(\omega_m).$$

In the case where $\mathrm{E}[Y]$ is unknown and sampling from Y is computationally less expensive than sampling from X, it is natural to estimate $\mathrm{E}[Y]$ using Monte Carlo sampling to yield a two-level Monte Carlo estimator of $\mathrm{E}[X]$ based on the control variate, Y, *i.e.*,

$$\tilde{\mu}_2 := \frac{1}{M_0}\sum_{m_0=1}^{M_0} Y(\omega_{m_0}) + \frac{1}{M_1}\sum_{m_1=1}^{M_1} (X-Y)(\omega_{m_1}).$$

See Section 1.6 for details about the definition of levels in our context.

In this work, we apply Giles's multilevel control variates idea to the hybrid Chernoff tau-leap approach to reduce the computational cost, which is measured as the amount of time needed for computing an estimate of $\mathrm{E}\left[g(X(T))\right]$, within TOL, with a given level of confidence. We show that our hybrid MLMC method has the same computational complexity of the pure SSA, *i.e.*, order $\mathcal{O}\left(TOL^{-2}\right)$. From this perspective, our method can be seen as a variance reduction for the SSA since our MLMC method does not change the complexity; it just reduces the corresponding multiplicative constant. We note in passing that in [2], the authors show that the computational complexity for the pure MLMC tau-leap case has order $\mathcal{O}\left(TOL^{-2}(\log(TOL))^2\right)$. We note also that here our goal is to provide an estimate of $\mathrm{E}\left[g(X(T))\right]$ in the probability sense and not in the mean square sense as in [1].

The global error arising from our hybrid tau-leap MLMC method can naturally be decomposed into three components: the global exit error, the time discretization error and the statistical error. This global error should be less than a prescribed tolerance, TOL, with probability larger than a certain confidence level. The global exit error is controlled by the one-step exit probability bound, δ [19]. The time discretization error, inherent to the tau-leap method, is controlled through the size of the mesh, Δt, [15]. At this point, it is crucial to stress that, by controlling the exit probability of the set of hybrid paths, we are indirectly turning this event into a rare event. Thus, direct sampling of exit paths is not an affordable way to estimate the probability of such an event.

Motivated by the Central Limit results of Collier et al. [6] for the Multilevel Monte Carlo estimator (see appendix A, Theorem 1), we approximate the statistical error with a Gaussian random variable with zero mean. In our numerical experiments, we tested this hypothesis by employing Q-Q plots and the Shapiro-Wilk test [20]. There, we did not reject the Gaussianity of the statistical error at the 1% significance level. The variance of the statistical error is a linear combination of the variance at the coarsest level and variances of the difference of two consecutive levels, which we sometimes call strong errors. In Section 3.3, motivated by the fact that sample variance and bias estimators are inaccurate on the deepest levels, we develop a novel dual-weighted residual expansion that allows us to estimate those quantities, cf. (3.3) and (3.4). We also control the statistical error through the number of coupled hybrid paths, $(M_\ell)_{\ell=0}^{L}$, simulated at each level.

We note that our use of duals in this work is different from the use in [15]. That earlier work proposed an adaptive, single-level, tau-leap algorithm for error control, choosing the time steps non-uniformly to control the global weak error based on dual-weighted error estimators. In this work, we do not have an adaptive time step based on dual-weighted error estimators as in [15]. We use instead dual-weighted error estimators to reduce the statistical error in our error estimates.

1.1. A Class of Markovian Pure Jump Processes. To describe the class of Markovian pure jump process, $X : [0, T] \times \Omega \to \mathbb{Z}_+^d$, that we use in this work, we consider a system of d species interacting through J different reaction channels. For the sake of brevity, we write $X(t, \omega) \equiv X(t)$. Let $X_i(t)$ be the number of particles of species i in the system at time t. We want to study the evolution of the state vector,

$$X(t) = (X_1(t), \ldots, X_d(t)) \in \mathbb{Z}_+^d,$$

modelled as a continuous-time, discrete-space Markov chain starting at some state, $X(0) \in \mathbb{Z}_+^d$. Each reaction can be described by the vector $\nu_j \in \mathbb{Z}^d$, such that, for a

state vector $x \in \mathbb{Z}_+^d$, a single firing of reaction j leads to the change

$$x \to x + \nu_j.$$

The probability that reaction j will occur during the small interval $(t, t+\mathrm{d}t)$ is then assumed to be

$$\mathrm{P}\left(X(t+\mathrm{d}t) = x + \nu_j | X(t) = x\right) = a_j(x)\mathrm{d}t + o\left(\mathrm{d}t\right), \tag{1.1}$$

with a given non-negative polynomial propensity function, $a_j : \mathbb{R}^d \to \mathbb{R}$. We set $a_j(x)$=0 for those x such that $x+\nu_j \notin \mathbb{Z}_+^d$. A process, X, that satisfies (1.1), is a continuous-time, discrete-space Markov chain that admits the following random time change representation [7]:

$$X(t) = X(0) + \sum_{j=1}^{J} \nu_j Y_j \left(\int_0^t a_j(X(s))\, \mathrm{d}s \right), \tag{1.2}$$

where $Y_j : \mathbb{R}_+ \times \Omega \to \mathbb{Z}_+$ are independent unit-rate Poisson processes. Hence, X is a non-homogeneous Poisson process.

In [15], the authors assume that there exists a vector, $w \in \mathbb{R}_+^d$, such that $(w, \nu_j) \leq 0$, for any reaction ν_j. Therefore, every reaction, ν_j, must have at least one negative component. This means that the species can be either transformed into other species or be consumed during the reaction. As a consequence, the space of states is contained in a simplex with vertices in the coordinate axis. This assumption excludes, for instance, birth processes. In our numerical examples, we allow the set of possible states of the system to be infinite, but we explicitly avoid cases in which one or more species grows exponentially fast or blows up in the time interval $[0, T]$.

REMARK 1.1. *In this setting, the solution of the following system of ordinary differential equations,*

$$\begin{cases} \dot{x}(t) &= \nu a(x(t)), \ t \in \mathbb{R}_+ \\ x(0) &= x_0 \in \mathbb{R}_+, \end{cases}$$

is called mean field solution, where ν is the matrix with columns ν_j and $a(x)$ is the column vector of propensities. In Section 4, we use the mean field path for scaling and preprocessing constants associated with the computational work of the SSA and Chernoff tau-leap steps.

1.2. Description of the Modified Next Reaction Method (MNRM). The MNRM, introduced in [3], based on the Next Reaction Method (NRM) [8], is an exact simulation algorithm like Gillespie's SSA that explicitly uses representation (1.2) for simulating exact paths and generates only one exponential random variable per iteration. The reaction times are modeled with firing times of Poisson processes, Y_j, with internal times given by the integrated propensity functions. The randomness is now separated from the state of the system and is encapsulated in the Y_j's. For each reaction, j, the internal time is defined as $R_j(t) = \int_0^t a_j(X(s))ds$. There are $J+1$ time frames in the system, the absolute one, t, and one for each Poisson process, Y_j. Computing the next reaction and its time is equivalent to computing how much time passes before one of the Poisson processes, Y_j, fires, and to identifying which process fires at that particular time, by taking the minimum of such times. The NRM and MNRM make use of internal times to reduce the number of simulated random

variables by half. In the following, we describe the MRNM and then we present its implementation in Algorithm 1.

Given t, we have the propensity $a_j{=}a_j(X(t))$ and the internal time $R_j{=}R_j(t)$. Now, let ΔR_j be the remaining time for the reaction, j, to fire, assuming that a_j stays constant over the interval $[t, t{+}\Delta R_j)$. Then, $t{+}\Delta R_j$ is the time when the next reaction, j, occurs. The next internal time at which the reaction, j, fires is then given by $R_j{+}a_j\Delta R_j$. When simulating the next step, the first reaction that fires occurs after $\Delta{=}\min_j \Delta R_j$. We then update the state of the system according to that reaction, add Δ to the global time, t, and then update the internal times by adding $a_j\Delta$ to each R_j. We are left to determine the value of ΔR_j, *i.e.*, the amount of time until the Poisson process, Y_j, fires, taking into account that a_j remains constant until the first reaction occurs. Denote by R_j the first firing time of Y_j that is strictly larger than R_j, *i.e.*, $P_j{:=}\min\{s{>}R_j : Y_j(s){>}Y_j(R_j)\}$ and finally $\Delta R_j{=}\frac{1}{a_j}(P_j - R_j)$.

Algorithm 1 The Modified Next Reaction Method. Inputs: the initial state, $X(0)$, the next grid point, $T_0 > t_0$, the propensity functions, $(a_j)_{j=1}^J$, the stoichiometric vectors, $(\nu_j)_{j=1}^J$. Outputs: the history of system states, $(X(t_k))_{k=0}^K$. Here, we denote $S \equiv (S_j)_{j=1}^J$, $P \equiv (P_j)_{j=1}^J$, and $R \equiv (R_j)_{j=1}^J$.

$k \leftarrow 0$, $t_k \leftarrow 0$, $X(t_k) \leftarrow X(0)$ and $R \leftarrow \mathbf{0}$
Generate J independent, uniform$(0,1)$ random numbers, r_j
$P \leftarrow (\log(1/r_j))_{j=1}^J$
while $t_k < T_0$ **do**
  $S \leftarrow (a_j(X(t_k)))_{j=1}^J$
  $(\Delta R_j)_{j=1}^J \leftarrow ((P_j - R_j)/S_j)_{j=1}^J$
  $\mu \leftarrow \text{argmin}_j\{\Delta R_j\}$
  $\Delta \leftarrow \min_j\{\Delta R_j\}$
  $t_{k+1} \leftarrow t_k + \Delta$
  $X(t_{k+1}) \leftarrow X(t_k) + \nu_\mu$
  $R \leftarrow R + S\Delta$
  $r \leftarrow \text{uniform}(0,1)$
  $P_\mu \leftarrow P_\mu + \log(1/r)$
  $k \leftarrow k{+}1$
end while
return $(X(t_l))_{l=0}^{k-1}$

Among the advantages already mentioned, we can easily modify Algorithm 1 to generate paths in the cases where the rate functions depend on time and also when there are reactions delayed in time. Finally, it is possible to simulate correlated exact/tau-leap paths using this algorithm as well as nested tau-leap/tau-leap paths. In [1], this technique is used to develop a uniform-step, unbiased, multilevel Monte Carlo (MLMC) algorithm. In Section 2.2, we use this feature for coupling two exact paths.

1.3. The Tau-Leap Approximation. In this section, we define $\bar{X}$, the tau-leap approximation of the process, X, which follows from applying the forward Euler approximation to the integral term in the following random time-change representa-

tion of X:

$$X(t+\tau) = X(t) + \sum_{j=1}^{J} \nu_j Y_j \left(\int_t^{t+\tau} a_j(X(s))\, \mathrm{d}s \right).$$

The tau-leap method was proposed in [11] to avoid the computational drawback of the exact methods, *i.e.*, when many reactions occur during a short time interval. The tau-leap process, $\bar{X}$, starts from $X(0)$ at time 0, and given that $\bar{X}(t)=\bar{x}$ and a time step $\tau>0$, we have that $\bar{X}$ at time $t+\tau$ is generated by

$$\bar{X}(t+\tau) = \bar{x} + \sum_{j=1}^{J} \nu_j \mathcal{P}_j \left(a_j(\bar{x})\tau \right),$$

where $\{\mathcal{P}_j(\lambda_j)\}_{j=1}^{J}$ are independent Poisson distributed random variables with parameter λ_j, used to model the number of times that the reaction j fires during the $(t, t+\tau)$ interval. Again, this is nothing other than a forward Euler discretization of the stochastic differential equation formulation of the pure jump process (1.2), realized by the Poisson random measure with state dependent intensity (see, *e.g.*, [17]).

In the limit, when τ tends to zero, the tau-leap method gives the same solution as the exact methods. The total number of firings in each channel is a Poisson-distributed stochastic variable depending only on the initial population, $\bar{X}(t)$. The error thus comes from the variation of $a(X(s))$ for $s \in (t, t+\tau)$.

We observe that the computational work of a tau-leap step involves the generation of J independent Poisson random variables. This is in contrast to the computational work of an exact step, which involves only the work of generating two uniform random variables, in the case of the SSA, and only one in the case of MNRM.

1.4. The Chernoff-Based Pre-Leap Check. In [19], we derived a Chernoff-type bound that allows us to guarantee that the one-step exit probability in the tau-leap method is less than a predefined quantity, $\delta>0$. We now briefly summarize the main idea. Consider the following pre-leap check problem: find the largest possible τ such that, with high probability, in the next step, the approximate process, $\bar{X}$, will take a value in the lattice, $\mathbb{Z}_+^d$, of non-negative integers. The solution to that problem can be achieved by solving d auxiliary problems, one for each x-coordinate, $i = 1, 2, \ldots, d$, as follows. Find the largest possible $\tau_i \geq 0$, such that

$$\mathrm{P}\left(\bar{X}_i(t) + \sum_{j=1}^{J} \nu_{ji} \mathcal{P}_j \left(a_j\left(\bar{X}(t)\right) \tau_i \right) < 0 \;\middle|\; \bar{X}(t) \right) \leq \delta_i, \tag{1.3}$$

where $\delta_i=\delta/d$, and ν_{ji} is the i-th coordinate of the j-th reaction channel, ν_j. Finally, we let $\tau := \min\{\tau_i : i = 1, 2, \ldots, d\}$. To find the largest time steps, τ_i, let $Q_i(t, \tau_i) := \sum_{j=1}^{J} (-\nu_{ji}) \mathcal{P}_j \left(a_j\left(\bar{X}(t)\right) \tau_i \right)$. Then, for all $s>0$, we have the Chernoff bound:

$$\mathrm{P}\left(Q_i(t,\tau_i) > \bar{X}_i(t) \,\middle|\, \bar{X}(t) \right) \leq \inf_{s>0} \exp\left(-s\bar{X}_i(t) + \tau_i \sum_{j=1}^{J} a_j(\bar{X}(t))(e^{-s\nu_{ji}} - 1) \right).$$

Expressing τ_i as a function of s, we write

$$\tau_i(s) = \frac{\log(\delta_i) + s\bar{X}_i(t)}{-a_0\left(\bar{X}(t)\right) + \sum_{j=1}^{J} a_j(\bar{X}(t)) e^{-s\,\nu_{ji}}} =: \frac{R_i(s)}{D_i(s)},$$

where

$$a_0(\bar{X}(t)) := \sum_{j=1}^{J} a_j(\bar{X}(t)).$$

We want to maximize τ_i while satisfying condition (1.3). Let τ_i^* be this maximum. We then have the following possibilities: If $\nu_{ji} \geq 0$, for all j, then naturally $\tau_i^* = +\infty$; otherwise, we have the following three cases:

1. $D_i(s_i)>0$. In this case, $\tau_i(s_i)=0$ and $D_i(s)$ is positive and increasing as $\forall s \geq s_i$. Therefore, $\tau_i(s)$ is equal to the ratio of two positive increasing functions. The numerator, $R_i(s)$, is a linear function and the denominator, $D_i(s)$, grows exponentially fast. Then, there exist an upper bound, τ_i^*, and a unique number, $\tilde{s}_i$, which satisfy $\tau_i(\tilde{s}_i)=\tau_i^*$. We developed an algorithm in [19] for approximating $\tilde{s}_i$, using the relation $\tau_i'(\tilde{s}_i)=0$.
2. If $D_i(s_i)<0$, then $\tau_i^* = +\infty$.
3. If $D_i(s_i)=0$, then $\tau_i^*=\bar{X}_i(t)/D_i'(s_i)$.

Here $s_i := -\log(\delta_i)/\bar{X}_i(t)$.

1.5. The Hybrid Algorithm for Single-Path Generation. In this section, we briefly summarize our previous work, presented in [19], on hybrid paths.

The main idea behind the hybrid algorithm is the following. A path generated by an exact method (like SSA or MNRM) never exits the lattice, $\mathbb{Z}_+^d$, although the computational cost may be unaffordable due to many small inter-arrival times typically occurring when the process is "far" from the boundary. A tau-leap path, which may be cheaper than an exact one, could leave the lattice at any step. The probability of this event depends on the size of the next time step and the current state of the approximate process, $\bar{X}(t)$. This one-step exit probability could be large, especially when the approximate process is "close" to the boundary. We developed in [19] a Chernoff-type bound to control the mentioned one-step exit probability. Even more, by construction, the probability that one hybrid path exits the lattice, $\mathbb{Z}_+^d$, can be estimated by

$$\mathrm{P}\left(A^c\right) \leq \mathrm{E}\left[1-(1-\delta)^{N_{\mathrm{TL}}}\right] = \delta \mathrm{E}\left[N_{\mathrm{TL}}\right] - \frac{\delta^2}{2}\left(\mathrm{E}\left[N_{\mathrm{TL}}^2\right] - \mathrm{E}\left[N_{\mathrm{TL}}\right]\right) + o(\delta^2),$$

where $\bar{\omega} \in A$ if and only if the whole hybrid path, $(\bar{X}(t_k,\bar{\omega}))_{k=0}^{K(\bar{\omega})}$, belongs to the lattice, $\mathbb{Z}_+^d$, $\delta>0$ is the one-step exit probability bound, and $N_{\mathrm{TL}}(\bar{\omega}) \equiv N_{\mathrm{TL}}$ is the number of tau-leap steps in a hybrid path. Here, A^c is the complement of the set A.

To simulate a hybrid exact/Chernoff tau-leap path, we first developed a one-step switching rule that, given the current state of the approximate process, $\bar{X}(t)$, adaptively determines whether to use an exact or an approximated method for the next step. This decision is based on the relative computational cost of taking an exact step (MNRM) versus the cost of taking a Chernoff tau-leap step. We show the switching rule in Algorithm 2. To compare the mentioned computational costs, we define K_1 as the ratio between the cost of computing τ_{Ch} and the cost of computing one step using the MNRM method, and $K_2=K_2(\bar{X}(t),\delta)$ is defined as the cost of taking a Chernoff tau-leap step, divided by the cost of taking a MNRM step plus the cost of computing τ_{Ch}. For further details on the switching rule, we refer to [19].

1.6. The Multilevel Monte Carlo Setting. In this subsection, we briefly summarize the control variates idea developed by Giles in [9]. Let $\{\bar{X}_\ell(t)\}_{t\in[0,T]}$ be

Algorithm 2 The one-step switching rule. Inputs: the current state of the approximate process, $\bar{X}(t)$, the current time, t, the values of the propensity functions evaluated at $\bar{X}(t)$, $(a_j(\bar{X}(t)))_{j=1}^J$, the one-step exit probability bound δ, and the next grid point, T_0. Outputs: method and τ. Notes: based on $\mathrm{E}\left[\tau_{\mathrm{SSA}}(\bar{X}(t)) \,\middle|\, \bar{X}(t)\right] = 1/a_0(\bar{X}(t))$ and $\tau_{Ch}(\bar{X}(t), \delta)$, this algorithm adaptively selects between MNRM and Chernoff tau-leap (TL). We denote by τ_{MNRM} (τ_{Ch}) the step size when the decision is to use the MNRM (tau-leap) method.

Require: $a_0 \leftarrow \sum_{j=1}^J a_j > 0$
if $K_1/a_0 < T_0 - t$ **then**
  $\tau_{Ch} \leftarrow$ compute Chernoff step size (see Section 2.2 in [19])
  if $\tau_{Ch} < K_2(\bar{X}(t), \delta)/a_0$ **then**
    return $(\mathrm{MNRM}, \tau_{\mathrm{MNRM}})$
  else
    return (TL, τ_{Ch})
  end if
else
  return $(\mathrm{MNRM}, \tau_{\mathrm{MNRM}})$
end if

a hybrid Chernoff tau-leap process with a time mesh of size Δt_ℓ and a one-step exit probability bound, δ. We can simulate paths of $\{\bar{X}_\ell(t)\}_{t\in[0,T]}$ by using Algorithm 4 in [19]. Let $g_\ell := g(\bar{X}_\ell(T))$.

Consider a hierarchy of nested meshes of the time interval $[0, T]$, indexed by $\ell = 0, 1, \ldots, L$. Let Δt_0 be the size of the coarsest time mesh that corresponds to the level ℓ=0. The size of the time mesh at level $\ell \geq 1$ is given by $\Delta t_\ell = R^{-\ell}\Delta t_0$, where R>1 is a given integer constant.

Assume that we are interested in estimating $\mathrm{E}[g_L]$, and we are able to simulate correlated pairs, $(g_\ell, g_{\ell-1})$ for $\ell = 1, \ldots, L$. Then, the following unbiased Monte Carlo estimator of $\mathrm{E}[g_L]$ uses g_{L-1} as a control variate:

$$\begin{aligned}\tilde{\mu}_L &:= \frac{1}{M_L}\sum_{m_L=1}^{M_L}\left(g_L(\omega_{m_L}) - (g_{L-1}(\omega_{m_L}) - \mathrm{E}[g_{L-1}])\right)\\ &= \mathrm{E}[g_{L-1}] + \frac{1}{M_L}\sum_{m_L=1}^{M_L}(g_L - g_{L-1})(\omega_{m_L}).\end{aligned}$$

Applying this idea recursively and taking into account the following telescopic decomposition: $\mathrm{E}[g_L] = \mathrm{E}[g_0] + \sum_{\ell=1}^L \mathrm{E}[g_\ell - g_{\ell-1}]$, we arrive at the multilevel Monte Carlo estimator of $\mathrm{E}[g_L]$:

$$\hat{\mu}_L := \frac{1}{M_0}\sum_{m_0=1}^{M_0} g_0(\omega_{m_0}) + \sum_{\ell=1}^{L}\frac{1}{M_\ell}\sum_{m_\ell=1}^{M_\ell}(g_\ell - g_{\ell-1})(\omega_{m_\ell}). \tag{1.4}$$

We have that $\hat{\mu}_L$ is unbiased, since $\mathrm{E}[\hat{\mu}_L] = \mathrm{E}[g_L]$. The variance of $\hat{\mu}_L$ is given by $\mathrm{Var}[\hat{\mu}_L] = \frac{\mathrm{Var}[g_0]}{M_0} + \sum_{\ell=1}^L \frac{\mathrm{Var}[g_\ell - g_{\ell-1}]}{M_\ell}$. Here, we are assuming independence among the batches between levels. For highly correlated pairs, $(g_\ell, g_{\ell-1})$, we can expect, for the same computational work, that $\mathrm{Var}[\hat{\mu}_L]$ is much less than the variance of the standard Monte Carlo estimator of $\mathrm{E}[g_L]$.

1.7. The Large Kurtosis Problem. Let us give a close examination of the problem of estimating $\mathrm{Var}\,[g_\ell - g_{\ell-1}]$ for highly correlated pairs, $(g_\ell, g_{\ell-1})$. This estimation is required to solve the optimization problem (4.3), that indicates how to choose the simulation parameters, particularly the number of simulated coupled paths for each pair of consecutive levels, $(M_\ell)_{\ell=0}^L$.

When ℓ becomes large, due to our coupling strategy developed in Section 2, we expect to obtain $g_\ell = g_{\ell-1}$ in most of our simulations, while observing differences only in a very small proportion of the simulated coupled paths.

For the sake of illustration, let us assume that the random variable $\chi_\ell := g_\ell - g_{\ell-1}$ takes values in the set $\{-1, 0, 1\}$, with respective probabilities $\{p_\ell, 1-2p_\ell, p_\ell\}$, where p_ℓ goes to zero. The kurtosis of χ_ℓ is by definition

$$\frac{\mathrm{E}\left[(\chi_\ell - \mathrm{E}\,[\chi_\ell])^4\right]}{\left(\mathrm{E}\left[(\chi_\ell - \mathrm{E}\,[\chi_\ell])^2\right]\right)^2} - 3.$$

Simple calculations show that the kurtosis of χ_ℓ is $(2p_\ell)^{-1}$, and we observe that $\chi_\ell^2 \sim$ Bernoulli$(2p_\ell)$. The maximum likelihood estimator of $2p_\ell$, $\hat{\theta}_\ell$, is the sample average of M_ℓ independent and identically distributed (iid) values of χ_ℓ^2. The coefficient of variation of $\hat{\theta}_\ell$, defined as $(\mathrm{Var}\left[\hat{\theta}_\ell\right])^{1/2}(\mathrm{E}\left[\hat{\theta}_\ell\right])^{-1}$, is $(2p_\ell M_\ell)^{-1}$. Therefore, an accurate estimation of p_ℓ requires a sample of size

$$M_\ell \gg (2p_\ell)^{-1} \to \infty.$$

This lower bound on M_ℓ goes strongly against the spirit of the Multilevel Monte Carlo method, where M_ℓ should be a decreasing function of ℓ.

To overcome this difficulty, in Section 3.3, we developed a formula based on dual-weighted residuals. The technique of dual-weighted residuals can be motivated as follows: consider a process $\bar{\bar{X}}$, such that its position at time s, having departed from the state x, at a previous time t, is denoted as $\bar{\bar{X}}(s;t,x)$. Notice that for $t<s<T$, we have that $\bar{\bar{X}}(T;t,x) = \bar{\bar{X}}(T;s,\bar{\bar{X}}(s;t,x))$. Let us define an auxiliary function $U(t,x) := g(\bar{\bar{X}}(T;t,x))$, where g is an observable scalar function of the final state of the process $\bar{\bar{X}}$ that started from the state x at the initial time, t. If $\bar{X}$ is a process approximating $\bar{\bar{X}}$, we want to have a computable approximation for $g(\bar{X}(T;0,x_0)) - g(\bar{\bar{X}}(T;0,x_0))$. Consider a time mesh, $\{0=t_0, t_1, \ldots, t_N=T\}$, and define $\bar{X}_{t_n} := \bar{X}(t_n;0,x_0)$, $\bar{X}_{t_{n+1}} := \bar{\bar{X}}(t_{n+1};t_n,\bar{X}_{t_n})$ and $e_{n+1} := \bar{X}_{t_{n+1}} - \bar{\bar{X}}_{t_{n+1}}$. Observe that

$$\begin{aligned}
g(\bar{X}(T;0,x_0)) - g(\bar{\bar{X}}(T;0,x_0)) &= U(T,\bar{X}(T;0,x_0)) - U(0,x_0) \\
&= \sum_{n=0}^{N-1} \left(U(t_{n+1},\bar{X}_{t_{n+1}}) - U(t_n,\bar{X}_{t_n})\right) \\
&= \sum_{n=0}^{N-1} \left(U(t_{n+1},\bar{X}_{t_{n+1}}) - U(t_{n+1},\bar{\bar{X}}(t_{n+1};t_n,\bar{X}_{t_n}))\right) \\
&= \sum_{n=0}^{N-1} \left(e_{n+1} \cdot \int_0^1 \nabla_x U(t_{n+1},\bar{X}_{t_{n+1}} - s e_{n+1}) ds\right) \\
&= \sum_{n=0}^{N-1} \left(e_{n+1} \cdot \nabla_x U(t_{n+1},\bar{X}_{t_{n+1}}) + \left\|\nabla^2 U\right\| \|e_{n+1}\|^2 + h.o.t.\right).
\end{aligned}$$

We can now write a backward recurrence for the dual weights, $(\phi_n)_{n=1}^N$:

$$\begin{aligned}
\phi_n :=& \nabla_x U(t_n, \bar{X}_{t_n}) = \partial_{\bar{X}_{t_n}} g(\bar{\bar{X}}(T; t_n, \bar{X}_{t_n})) \\
&= \partial_{\bar{X}_{t_n}} g(\bar{\bar{X}}(T; t_{n+1}, \bar{X}_{t_{n+1}})) \\
&= \partial_{\bar{X}_{t_{n+1}}} g(\bar{\bar{X}}(T; t_{n+1}, \bar{X}_{t_{n+1}})) \frac{\partial_{\bar{X}_{t_{n+1}}}}{\partial_{\bar{X}_{t_n}}} \\
&= \nabla_x U(t_{n+1}, \bar{X}_{t_{n+1}}) \frac{\partial_{\bar{X}_{t_{n+1}}}}{\partial_{\bar{X}_{t_n}}} \\
&= \phi_{n+1} \frac{\partial_{\bar{X}_{t_{n+1}}}}{\partial_{\bar{X}_{t_n}}} \\
\phi_N :=& \nabla g(\bar{X}(T; 0, x_0)).
\end{aligned}$$

This reasoning evidently works for processes $\bar{\bar{X}}$ that are pathwise differentiable with respect to the initial condition. Our space state is in general a subset of the lattice, $\mathbb{Z}_+^d$, and for that reason, we can not directly apply this technique. In [14], the authors show how this dual-weighted residual technique can be adapted to the tau-leap case in regimes close to the mean field or to the Stochastic Langevin limit. In more general regimes, the formula (3.4), which provides accurate estimates of $\mathrm{Var}\,[g_\ell - g_{\ell-1}]$ in our numerical examples (see for instance Figure 5.5 in Section 5), is promising but more research is needed in this direction. Specifically, in Section 3.3, the formula (3.4) is deduced from the conditional distribution of the local errors, $e_{n+1}|\mathcal{F}$, conditional on a sigma-algebra, $\mathcal{F}$, generated by the sequence, $(\bar{X}_{t_n})_{n=1}^N$, and applying the tower properties of conditional expectation and conditional variance. Similar comments apply to Formula (3.3) regarding the weak error, $\mathrm{E}\,[g(X(T)) - g_L]$.

1.8. Outline of this Work. In Section 2, we first show the main idea for coupling two tau-leap paths, which comes from a construction by Kurtz [16] for coupling two Poisson random variables. Then, inspired by the ideas of Anderson and Higham in [1], we propose an algorithm for coupling two hybrid Chernoff tau-leap paths (see [19]). This algorithm uses four building blocks that result from the combination of the MNRM and the tau-leap methods. In Section 3, we propose a novel hybrid MLMC estimator. Next, we introduce a global error decomposition; and finally, we develop formulae to efficiently estimate the variance of the difference of two consecutive levels and to estimate the bias based on dual-weighted residuals. These estimates are particularly useful to addressing the large kurtosis problem, described in Section 1.7, that appears at the deeper levels and makes standard sample estimators too costly. Next, in Section 4, we show how to control the three error components of the global error and how to obtain the parameters needed for computing the hybrid MLMC estimator to achieve a given tolerance with nearly optimal computational work. We also show that the computational complexity of our method is of order $\mathcal{O}\left(TOL^{-2}\right)$. In Section 5, the numerical examples illustrate the advantages of the hybrid MLMC method over the single-level approach presented in [19] and to the SSA. Section 6 presents our conclusions and suggestions for future work.

2. Generating Coupled Hybrid Paths. In this section, we present an algorithm that generates coupled hybrid Chernoff tau-leap paths, which is an essential ingredient for the multilevel Monte Carlo estimator. We first show how to couple two Poisson random variables and then we explain how we make use of the two algorithms

presented in [1] as Algorithms 2 and 3 and two additional algorithms we developed to create an algorithm that generates coupled hybrid paths.

2.1. Coupling Two Poisson Random Variables. We motivate our coupling algorithm (Algorithm 3) by first describing how to couple two Poisson random variables. In our context, 'coupling' means that we want to induce a correlation between them that is as strong as possible. This construction was first proposed by Kurtz in [16]. Suppose that we want to couple $\mathcal{P}_1(\lambda_1)$ and $\mathcal{P}_2(\lambda_2)$, two Poisson random variables, with rates λ_1 and λ_2, respectively. Consider the following decompositions,

$$\begin{aligned}\mathcal{P}_1(\lambda_1) &:= \mathcal{P}^*(\lambda_1 \wedge \lambda_2) + \mathcal{Q}_1(\lambda_1 - \lambda_1 \wedge \lambda_2)\\ \mathcal{P}_2(\lambda_2) &:= \mathcal{P}^*(\lambda_1 \wedge \lambda_2) + \mathcal{Q}_2(\lambda_2 - \lambda_1 \wedge \lambda_2),\end{aligned}$$

where $\mathcal{P}^*(\lambda_1 \wedge \lambda_2)$, $\mathcal{Q}_1(\lambda_1 - \lambda_1 \wedge \lambda_2)$ and $\mathcal{Q}_2(\lambda_2 - \lambda_1 \wedge \lambda_2)$ are three independent Poisson random variables. Here, $\lambda_1 \wedge \lambda_2 := \min\{\lambda_1, \lambda_2\}$. Observe that at least one of the following vanishes: $\mathcal{Q}_1(\lambda_1 - \lambda_1 \wedge \lambda_2)$ and $\mathcal{Q}_2(\lambda_2 - \lambda_1 \wedge \lambda_2)$. This is because at least one of the rates is zero. Algorithm 3 implements these ideas. Finally, note that, by construction, we have

$$\begin{aligned}\mathrm{Var}\left[\mathcal{P}_1(\lambda_1) - \mathcal{P}_2(\lambda_2)\right] &= \mathrm{Var}\left[\mathcal{Q}_1(\lambda_1 - \lambda_1 \wedge \lambda_2) - \mathcal{Q}_2(\lambda_2 - \lambda_1 \wedge \lambda_2)\right]\\ &= \left|\lambda_1 - \lambda_2\right|.\end{aligned}$$

However, if instead we consider making $\mathcal{P}_1(\lambda_1)$ and $\mathcal{P}_2(\lambda_2)$ independent, then

$$\mathrm{Var}\left[\mathcal{P}_1(\lambda_1) - \mathcal{P}_2(\lambda_2)\right] = \lambda_1 + \lambda_2,$$

which may be a large value even when λ_1 and λ_2 are close.

2.2. Coupling Two Hybrid Paths. In this section, we describe how to generate two coupled hybrid Chernoff tau-leap paths, $\bar{X}$ and $\bar{\bar{X}}$, corresponding to two nested time discretizations, called coarse and fine, respectively. Assume that the current time is t, and we know the states, $\bar{X}(t)$ and $\bar{\bar{X}}(t)$. Based on this knowledge, we have to determine a method for each level. This method can be either the MNRM or the tau-leap one, determining four possible combinations leading to four algorithms, B1, B2, B3 and B4, that we use as building blocks. Table 2.1 summarizes them.

	Algorithm	at coarser mesh	at fine mesh
B1	(part of Algorithm 3)	TL	TL
B2	(Algorithm 5)	TL	MNRM
B3	"	MNRM	TL
B4	"	MNRM	MNRM

TABLE 2.1

Building blocks for simulating two coupled hybrid Chernoff tau-leap paths. Algorithms B1 and B2 are presented as Algorithms 2 and 3 in [1]. Algorithm B3 can be directly obtained from Algorithm B2. Algorithm B4 is also based on Algorithm B2, but to produce MNRM steps, we update the propensities at the coarse level at the beginning of each time interval defined by the fine level.

We note that the only case in which we use a Poisson random variates generator for the tau-leap method is in Algorithm B1. In Algorithms B2 and B3, the Poisson random variables are simulated by adding independent exponential random variables

with the same rate, λ, until a given time final time T is exceeded. The rate, λ, is obtained by freezing the propensity functions, a, at time t. More specifically, the Poisson random variates are obtained by using the MNRM repeatedly without updating the intensity.

We now briefly describe the Chernoff hybrid coupling algorithm, *i.e.*, Algorithm 3. Given the current time, t, and the current state of the process at the coarse level, $\bar{X}(t)$, and the fine level, $\bar{\bar{X}}(t)$, this algorithm determines the next time point at which we run the algorithm (called time "horizon"). To fix the idea, let us assume that, based on $\bar{X}(t)$, the one-step switching rule, *i.e.*, Algorithm 2, chooses the tau-leap method at the coarse level, with the corresponding Chernoff step size, $\bar{\tau}$. As we mentioned, this $\bar{\tau}$ is the largest step size such that the probability that the process, in the next time step, takes a value outside $\mathbb{Z}_+^d$, is less than $\bar{\delta}$. This step size plus the current time, t, cannot be greater than the final time, T, and also cannot be greater than the next time discretization grid point in the coarse grid, $\bar{t}$, because the discretization error must be controlled. Taking the minimum of all those values, we obtain the next time horizon at the coarse grid, $\bar{H}$. Note that, if the chosen method is MNRM instead of tau-leap, we do not need to take into account the grid, and the next time horizon will be the minimum between the next reaction time and the final time, T.

We now explain algorithm B1 (TL-TL). Assume that tau-leap is chosen at the coarse and at the fine level. We thus obtain two time horizons, one for the coarse level, $\bar{H}$, and another for the fine level, $\bar{\bar{H}}$. In this case, the global time horizon will be $H{:=}\min\{\bar{H},\bar{\bar{H}}\}$. Since the chosen method in both grid levels is tau-leap, we need to freeze the propensities at the beginning of the corresponding intervals. In the coarse case, during the interval $[t,\bar{H})$ (the propensities are equal to $a(\bar{X}(t)){=:}\bar{a}$), and in the fine case during the interval $[t,\bar{\bar{H}})$ (the propensities are equal to $a(\bar{\bar{X}}(t)){=:}\bar{\bar{a}}$). Suppose that $\bar{H}<\bar{\bar{H}}$ (see Figure 2.1).

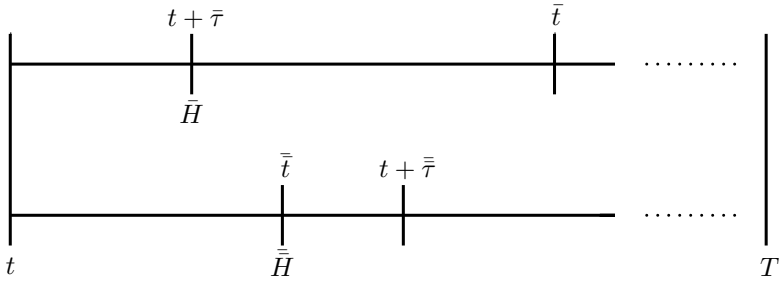

FIG. 2.1. *This figure depicts a particular instance of the Chernoff hybrid coupling algorithm (Algorithm 3), where $\bar{\tau}<\bar{\bar{\tau}}$. The synchronization horizon H, defined as $H{:=}\min\{\bar{H},\bar{\bar{H}}\}$, is equal to $\bar{H}$ in this case. Notice that $\bar{H}{:=}\min\{\bar{t},t+\bar{\tau},T\}$ and $\bar{\bar{H}}{:=}\min\{\bar{\bar{t}},t+\bar{\bar{\tau}},T\}$*

Then, we couple two Poisson random variables at time $t{=}\bar{H}$, using the idea described in Section 2.1. When time reaches $\bar{H}$, the decision between which method to use (and the corresponding step size) at the coarse level must be made again. Note that the propensities of the process at the fine grid will be kept frozen until $\bar{\bar{H}}$. The case when $\bar{H}>\bar{\bar{H}}$ is analogous to the one we described, but the decisions on the method and step size are made at the finer level, when time reaches $\bar{\bar{H}}$. It can also be possible that $\bar{H}=\bar{\bar{H}}$. In that case, the decision between which method to use (and the corresponding step size) must be made at the coarse and at the fine level.

In the case of algorithm B2 (TL-MNRM), we assume that tau-leap is chosen at the coarse level, and MNRM at the fine level, obtaining two time horizons, one for the coarse level, $\bar{H}$, and another for the fine level, $\bar{\bar{H}}$. The only difference in how we determine the time horizons between algorithms B1 and B2 is that the time discretization grid points in the fine grid are not taken into account to determine $\bar{\bar{H}}$. Algorithm B2 is then applied until the simulation reaches $H := \min\{\bar{H}, \bar{\bar{H}}\}$. Suppose that $\bar{\bar{H}} < \bar{H}$. In this case, the process $\bar{\bar{X}}$ could take more than one step to reach $\bar{\bar{H}}$. At each step, the propensity functions $a(\bar{\bar{X}}(\cdot))$ are computed, but not the propensities for the coarse level, because in that case the tau-leap method is used. Note that the decision between which algorithm to use (B2 or another) is not made at those steps, but only when time reaches $\bar{\bar{H}}$. When time reaches $\bar{\bar{H}}$, the decision of which method to use (and the corresponding step size) at the fine level must be made again. In this case, the propensities at the coarse grid will be kept frozen until $\bar{H}$. The reasoning for the cases $\bar{\bar{H}} > \bar{H}$ and $\bar{\bar{H}} = \bar{H}$ are similar to before.

The other two cases, that is, B3 and B4, are the same as B2. The only difference resides is when to update the propensity values, $\bar{a}$ and $\bar{\bar{a}}$. See Algorithm 3 for more details. As made clear in the preceding paragraphs, the decision on which algorithm to use for a certain time interval is made only at the horizon points.

REMARK 2.1. *[About telescoping] To ensure the telescoping sum property, the probability law of the hybrid process at level ℓ should be the same disregarding whether level ℓ is the finer in the pair $(\bar{X}_{\ell-1}, \bar{\bar{X}}_{\ell})$ or the coarser in the pair $(\bar{X}_{\ell}, \bar{\bar{X}}_{\ell+1})$. For that reason, each process has its own next horizon as its decision points. See Figure 2.1 showing the time horizons scheme and Figures 5.14 and 5.15 in Section 5 to see that the telescoping sum property is satisfied by our hybrid coupling sampling scheme.*

3. Multilevel Monte Carlo Estimator and Global Error Decomposition.

In this section, we present the multilevel Monte Carlo estimator. We first show the estimator and its properties and then we analyze and control the computational global error, which is decomposed into three error components: the discretization error, the global exit error, and the Monte Carlo statistical error. We give upper bounds for each one of the three components.

3.1. The MLMC Estimator.

In this section, we discuss and implement a variation of the multilevel Monte Carlo estimator (1.4) for the hybrid Chernoff tau-leap case. The main ingredient of this section is Algorithm 3, which generates coupled hybrid paths at levels $\ell-1$ and ℓ. Let us now introduce some notation. Let A_ℓ be the event in which the $\bar{X}_\ell$-path arrived at the final time, T, without exiting the state space of X. Let $\mathbf{1}_A$, be the indicator function of an arbitrary set, A. Finally, $g_\ell := g(\bar{X}_\ell(T))$ was defined in Section 1.6.

Consider the following telescopic decomposition:

$$\mathrm{E}\left[g_L \mathbf{1}_{A_L}\right] = \mathrm{E}\left[g_0 \mathbf{1}_{A_0}\right] + \sum_{\ell=1}^{L} \mathrm{E}\left[g_\ell \mathbf{1}_{A_\ell} - g_{\ell-1}\mathbf{1}_{A_{\ell-1}}\right],$$

which motivates the definition of our MLMC estimator of $\mathrm{E}\left[g(X(T))\right]$,

$$\mathcal{M}_L := \frac{1}{M_0}\sum_{m=1}^{M_0} g_0 \mathbf{1}_{A_0}(\omega_{m,0}) + \sum_{\ell=1}^{L}\frac{1}{M_\ell}\sum_{m=1}^{M_\ell}[g_\ell \mathbf{1}_{A_\ell} - g_{\ell-1}\mathbf{1}_{A_{\ell-1}}](\omega_{m,\ell}). \tag{3.1}$$

3.2. Global Error Decomposition. In this section, we define the computational global error, $\mathcal{E}_L$, and show how it can be naturally decomposed into three components: the discretization error, $\mathcal{E}_{I,L}$, and the exit error, $\mathcal{E}_{E,L}$, both coming from the tau-leap part of the hybrid method and the Monte Carlo statistical error, $\mathcal{E}_{S,L}$. Next, we show how to model and control the global error, $\mathcal{E}_L$, giving upper bounds for each one of the three components. We define the computational global error, $\mathcal{E}_L$, as

$$\mathcal{E}_L := \mathrm{E}\left[g(X(T))\right] - \mathcal{M}_L.$$

Now, consider the following decomposition of $\mathcal{E}_L$:

$$\begin{aligned}\mathrm{E}\left[g(X(T))\right] - \mathcal{M}_L &= \mathrm{E}\left[g(X(T))(\mathbf{1}_{A_L} + \mathbf{1}_{A_L^c})\right] - \mathrm{E}\left[g_L \mathbf{1}_{A_L}\right] + \mathrm{E}\left[g_L \mathbf{1}_{A_L}\right] - \mathcal{M}_L \\ &= \underbrace{\mathrm{E}\left[g(X(T))\mathbf{1}_{A_L^c}\right]}_{=:\mathcal{E}_{E,L}} + \underbrace{\mathrm{E}\left[(g(X(T)) - g_L)\,\mathbf{1}_{A_L}\right]}_{=:\mathcal{E}_{I,L}} + \underbrace{\mathrm{E}\left[g_L \mathbf{1}_{A_L}\right] - \mathcal{M}_L}_{=:\mathcal{E}_{S,L}}.\end{aligned}$$

We show in [19] that by choosing adequately the one-step exit probability bound, δ, the exit error, $\mathcal{E}_{E,L}$, satisfies $|\mathcal{E}_{E,L}| \leq |\mathrm{E}\left[g(X(T))\right]|\,\mathrm{P}\left(A_L^c\right) \leq TOL^2$. An efficient procedure for accurately estimating $\mathcal{E}_{I,L}$ in the context of the tau-leap method is described in [15]. We adapt this method in Algorithm 9 for estimating the weak error in the hybrid context. A brief description follows. For each hybrid path, $(\bar{X}_\ell(t_{n,\ell},\bar{\omega}))_{n=0}^{N(\bar{\omega})}$, we define the sequence of dual weights $(\varphi_{n,\ell}(\bar{\omega}))_{n=1}^{N(\bar{\omega})}$ backwards as follows (see Section 1.7):

$$\begin{aligned}\varphi_{N(\bar{\omega}),\ell} &:= \nabla g(\bar{X}_\ell(t_{N(\bar{\omega}),\ell},\bar{\omega})) \\ \varphi_{n,\ell} &:= \left(Id + \Delta t_{n,\ell}\mathbb{J}_a^T(\bar{X}_\ell(t_{n,\ell},\bar{\omega}))\,\nu^T\right)\varphi_{n+1,\ell}, \quad n = N(\bar{\omega})-1,\ldots,1,\end{aligned} \tag{3.2}$$

where $\Delta t_{n,\ell} := t_{n+1,\ell} - t_{n,\ell}$, ∇ is the gradient operator and $\mathbb{J}_a(\bar{X}_\ell(t_{n,\ell},\bar{\omega})) \equiv [\partial_i a_j(\bar{X}_\ell(t_{n,\ell},\bar{\omega}))]_{j,i}$ is the Jacobian matrix of the propensity function, a_j, for $j=1\ldots J$ and $i=1\ldots d$. According to this method, $\mathcal{E}_{I,L}$ is approximated by $\mathcal{A}\left(\mathcal{E}_{I,L}(\bar{\omega});\cdot\right)$, where

$$\mathcal{E}_{I,L}(\bar{\omega}) := \sum_{n=1}^{N(\bar{\omega})} \left(\frac{\Delta t_{n,L}}{2}\mathbf{1}_{TL}(n)\sum_{j=1}^{J}(\varphi_{n,L}\cdot\nu_j)\Delta a_{j,n}\right)(\bar{\omega}), \tag{3.3}$$

$\mathcal{A}\left(X;M\right) := \frac{1}{M}\sum_{m=1}^{M} X(\omega_m)$, and, $\mathcal{S}^2\left(X;M\right) := \mathcal{A}\left(X^2;M\right) - \mathcal{A}\left(X;M\right)^2$ denote the sample mean and the sample variance of the random variable, X, respectively. Here, $\Delta a_{j,n}(\bar{\omega}) := a_j(\bar{X}_L(t_{n+1,\ell},\bar{\omega})) - a_j(\bar{X}_L(t_{n,\ell},\bar{\omega}))$, $\mathbf{1}_{TL}(n)=1$ if and only if, at time $t_{n,\ell}$, the tau-leap method was used, and we denote by Id the $d\times d$ identity matrix.

REMARK 3.1 (Computational cost of dual computations). *It is easy to see that the computational cost per path of the dual computations in (3.2) is comparable, and possibly smaller than the hybrid path. Indeed, no new random variables, especially Poisson ones, which are the most computationally expensive in the forward simulation, need to be sampled and no coupling between levels is needed. Moreover, we use (3.2) only to determine the discretisation parameters for the actual run; so (3.2) is thus used only in a fraction of the realisations.*

The variance of the statistical error, $\mathcal{E}_{S,L}$, is given by $\sum_{\ell=0}^{L}\frac{\mathcal{V}_\ell}{M_\ell}$, where $\mathcal{V}_0 := \mathrm{Var}\left[g_0\mathbf{1}_{A_0}\right]$ and $\mathcal{V}_\ell := \mathrm{Var}\left[g_\ell\mathbf{1}_{A_\ell} - g_{\ell-1}\mathbf{1}_{A_{\ell-1}}\right]$, $\ell \geq 1$. In the next subsection, we show how to estimate $\mathcal{V}_\ell$ efficiently using the duals from (3.2).

3.3. Dual-weighted Residual Estimation of $\mathcal{V}_\ell$. Here, we derive the formula (3.4) for estimating the variance, $\mathcal{V}_\ell$, $\ell \geq 1$. It is based on dual-weighted local errors arising from two consecutive tau-leap approximations of the process, X. For each level $\ell \geq 1$, the formula estimates $\mathcal{V}_\ell$ with much smaller statistical error than the standard sample estimator, which is seriously affected by the large kurtosis present at the deepest levels (see Section 1.7).

Let us introduce some notation:

$$
\begin{aligned}
f_{j,n} &:= (\varphi_{n+1} \cdot \nu_j),\\
\mu_{j,n} &:= \frac{\Delta t_n}{2} \sum_i (\nabla a_j(x_n) \cdot \nu_i) a_i(x_n),\\
\bar{\mu}_{j,n} &:= \frac{\Delta t_n}{2} \sum_i |(\nabla a_j(x_n) \cdot \nu_i)| a_i(x_n),\\
\sigma^2_{j,n} &:= \frac{\Delta t_n}{2} \sum_i (\nabla a_j(x_n) \cdot \nu_i)^2 a_i(x_n),\\
m_{j,n} &:= \min\{\bar{\mu}_{j,n}, \sqrt{\mu^2_{j,n} + \sigma^2_{j,n}}\},\\
q_{j,n} &:= \frac{\mu_{j,n}}{\sigma_{j,n}},\\
p_{j,n} &:= \Phi(-q_{j,n}),\\
\tilde{\mu}_{j,n} &:= \mu_{j,n}(1-2p_{j,n}),\\
\tilde{\sigma}_{j,n} &:= \sqrt{\frac{2}{\pi}} \sigma_{j,n} \exp(-q^2_{j,n}/2).
\end{aligned}
$$

Here, $\Phi(x)$ is the cumulative distribution function of a standard Gaussian random variable. We define our dual-weighted estimator of $\mathcal{V}_\ell$ as

(3.4)

$$
\begin{aligned}
\hat{\mathcal{V}}_\ell := \mathcal{S}^2 &\left(\sum_n \mathbf{1}_{TL}(n) \frac{\Delta t_n}{2} \sum_j f_{j,n} \mu_{j,n}; M_\ell \right)\\
&+ \mathcal{A}\left(\sum_n \mathbf{1}_{TL}(n) \frac{(\Delta t_n)^3}{8} \sum_{j,j'} f_{j,n} f_{j',n} \sum_i (\nabla a_j(x_n) \cdot \nu_i)(\nabla a_{j'}(x_n) \cdot \nu_i) a_i(x_n); M_\ell \right)\\
&+ \mathcal{A}\left(\sum_n \mathbf{1}_{TL}(n) \frac{\Delta t_n}{2} \sum_j f^2_{j,n} \left(\mathbf{1}_{G_n} (\tilde{\mu}_{j,n} + \tilde{\sigma}_{j,n}) + \mathbf{1}_{G_n^c} m_{j,n} \right), M_\ell \right),
\end{aligned}
$$

where $\mathbf{1}_{G_n}$=1 if and only if $a_j(x_n)\frac{\Delta t_n}{2} > c$ for all $j \in \{1, \ldots, J\}$, where c is a positive user-defined constant.

First, notice that $\mathcal{V}_\ell$ could be a very small positive number. In fact, in our numerical experiments, we observe that the standard Monte Carlo sample estimation of this quantity turns out to be computationally infeasible due to the huge number of simulations required to stabilize its coefficient of variation. For this reason, we

initially consider the following dual-weighted approximations:

$$\mathrm{E}\,[\,g_\ell - g_{\ell-1}] \approx \mathrm{E}\left[\sum_n \varphi_{n+1,\ell-1}\cdot e_{n+1,\ell-1}\right], \tag{3.5}$$

$$\mathrm{Var}\,[\,g_\ell - g_{\ell-1}] \approx \mathrm{Var}\left[\sum_n \varphi_{n+1,\ell-1}\cdot e_{n+1,\ell-1}\right],$$

where $(\varphi_{n+1,\ell-1})_{n=0}^{N(\bar{\omega})-1}$, defined in (3.2), is a sequence of dual weights computed backwards from a simulated path, $(\bar{X}_\ell(t_{n,\ell-1}))_{n=1}^{N(\bar{\omega})}$, and the sequence of local errors, $(e_{n+1,\ell-1})_{n=0}^{N(\bar{\omega})-1}$, defined in (3.10), is the subject of the next subsection.

Defining the Sequence of Local Errors. For simplicity of analysis, we make two assumptions: i) the time mesh associated with the level, ℓ, is obtained by halving the intervals of the level $\ell-1$; ii) we perform the tau-leap at both levels without considering the Chernoff bounds described in Section 1.4.

Let $\bar{X}$ and $\bar{\bar{X}}$ be two tau-leap approximations of X based on two consecutive grid levels, for instance, $\bar{X}{:=}\bar{X}_{\ell-1}$ and $\bar{\bar{X}}{:=}\bar{X}_\ell$. Consider two consecutive time-mesh points for $\bar{X}$, $\{t_n, t_{n+1}\}$, and three consecutive time-mesh points for $\bar{\bar{X}}$, $\{t_n, (t_n{+}t_{n+1})/2, t_{n+1}\}$. Let $\bar{X}$ and $\bar{\bar{X}}$ start from x_n at time t_n.

The first step for coupling $\bar{X}$ and $\bar{\bar{X}}$ is to define

$$\bar{X}_{n+1} := x_n + \sum_j \nu_j \mathcal{Y}_{j,n}(a_j(x_n)\Delta t_n), \tag{3.6}$$

$$Z_{n+1} := x_n + \sum_j \nu_j \mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}), \tag{3.7}$$

$$\bar{\bar{X}}_{n+1} := Z_{n+1} + \sum_j \nu_j \mathcal{R}_{j,n}(a_j(Z_{n+1})\frac{\Delta t_n}{2}),$$

where $\{\mathcal{Y}_{j,n}\}_{j=1}^J \cup \{\mathcal{Q}_{j,n}\}_{j=1}^J \cup \{\mathcal{R}_{j,n}\}_{j=1}^J$ are Poisson random variables. To couple the $\bar{X}$ and $\bar{\bar{X}}$ processes, we first decompose $\mathcal{Y}_{j,n}(a_j(x_n)\Delta t_n)$ as the sum of two independent Poisson random variables, $\mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) + \mathcal{Q}'_{j,n}(a_j(x_n)\frac{\Delta t_n}{2})$. As a consequence, $\bar{X}$ and $\bar{\bar{X}}$ coincide in the closed interval $[t_n, (t_n{+}t_{n+1})/2]$. By applying this decomposition in (3.6), we obtain

$$\bar{X}_{n+1} = x_n + \sum_j \nu_j \mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) + \sum_j \nu_j \mathcal{Q}'_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}), \tag{3.8}$$

$$\bar{\bar{X}}_{n+1} = x_n + \sum_j \nu_j \mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) + \sum_j \nu_j \mathcal{R}_{j,n}(a_j(Z_{n+1})\frac{\Delta t_n}{2}).$$

The second step for coupling $\bar{X}$ and $\bar{\bar{X}}$, according to [1], is as follows: let $m_j := \min\{a_j(x_n), a_j(Z_{n+1})\}$, $c_j := a_j(x_n) - m_j$ and $f_j := a_j(Z_{n+1}) - m_j$. Notice that for each j, either c_j or f_j is zero (or both).

Now, consider the following decompositions:

$$\mathcal{Q}'_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) = \mathcal{P}'_{j,n}(m_j\frac{\Delta t_n}{2}) + \mathcal{P}''_{j,n}(c_j\frac{\Delta t_n}{2}), \tag{3.9}$$

$$\mathcal{R}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) = \mathcal{P}'_{j,n}(m_j\frac{\Delta t_n}{2}) + \mathcal{R}'_{j,n}(f_j\frac{\Delta t_n}{2}),$$

where $\mathcal{P}'$, $\mathcal{P}''$ and $\mathcal{R}'$ are independent Poisson random variables.

By substituting (3.9) into (3.8), we define the local error, $e_{n+1,\ell-1}$, as

(3.10)
$$\begin{aligned} e_{n+1,\ell-1} &:= \bar{\bar{X}}_{n+1} - \bar{X}_{n+1} \\ &= \sum_j \nu_j \left(\mathcal{R}'_{j,n}(f_j \frac{\Delta t_n}{2}) - \mathcal{P}''_{j,n}(c_j \frac{\Delta t_n}{2}) \right) \\ &= \sum_j \nu_j \left(\mathcal{R}'_{j,n}(\Delta a_{j,n} \frac{\Delta t_n}{2}) \mathbf{1}_{\{\Delta a_{j,n} > 0\}} - \mathcal{P}''_{j,n}(-\Delta a_{j,n} \frac{\Delta t_n}{2}) \mathbf{1}_{\{\Delta a_{j,n} < 0\}} \right), \end{aligned}$$

where $\Delta a_{j,n} := a_j(Z_{n+1}) - a_j(x_n)$ and Z_{n+1} is defined in (3.7). Note that in (3.10) not only are $\mathcal{R}'_{j,n}$ and $\mathcal{P}''_{j,n}$ random variables, but $\Delta a_{j,n}$ is also random because it depends on the random variables $(\mathcal{Q}_{j,n})_{j=1}^J$. Also note that all the mentioned random variables are independent.

Conditioning. At this moment, it is convenient to recall the tower properties of the conditional expectation and the conditional variance: given a random variable, X, and a sigma algebra, $\mathcal{F}$, defined over the same probability space, we have

(3.11)
$$\begin{aligned} \mathrm{E}\left[X\right] &= \mathrm{E}\left[\mathrm{E}\left[X \,\middle|\, \mathcal{F}\right]\right], \\ \mathrm{Var}\left[X\right] &= \mathrm{Var}\left[\mathrm{E}\left[X \,\middle|\, \mathcal{F}\right]\right] + \mathrm{E}\left[\mathrm{Var}\left[X \,\middle|\, \mathcal{F}\right]\right]. \end{aligned}$$

Hereafter, we fix ℓ and, for the sake of brevity, omit it as a subindex.

Applying (3.11) to $\sum_n \varphi_{n+1} \cdot e_{n+1}$ and conditioning on $\mathcal{F}$, we obtain

$$\begin{aligned} \mathrm{Var}\left[\sum_n \varphi_{n+1} \cdot e_{n+1}\right] &= \mathrm{Var}\left[\mathrm{E}\left[\sum_n \varphi_{n+1} \cdot e_{n+1} \middle| \mathcal{F}\right]\right] + \mathrm{E}\left[\mathrm{Var}\left[\sum_n \varphi_{n+1} \cdot e_{n+1} \middle| \mathcal{F}\right]\right] \\ &= \mathrm{Var}\left[\sum_n \mathrm{E}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right]\right] + \mathrm{E}\left[\sum_n \mathrm{Var}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right]\right]. \end{aligned}$$

The main idea is to generate M_ℓ Monte Carlo paths, $(\bar{X}_\ell(t_n;\bar{\omega}))_{n=1}^{N(\bar{\omega})}$, and to estimate $\mathrm{Var}\left[\sum_n \varphi_{n+1} \cdot e_{n+1}\right]$ using

(3.12)
$$\hat{\mathcal{V}}_\ell := \mathcal{S}^2 \left(\underbrace{\sum_n \mathrm{E}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right](\bar{\omega})}_{S_e(\bar{\omega})}; M_\ell \right) + \mathcal{A} \left(\underbrace{\sum_n \mathrm{Var}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right](\bar{\omega})}_{S_v(\bar{\omega})}; M_\ell \right).$$

To avoid nested Monte Carlo calculations, we develop exact and approximate formulas for computing $\mathrm{E}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right]$ and $\mathrm{Var}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right]$. To derive those formulas, we consider a sigma-algebra, $\mathcal{F}$, such that $(\varphi_n(\bar{\omega}))_{n=1}^{N(\bar{\omega})}$, conditioned on $\mathcal{F}$, is deterministic, *i.e.*, $(\varphi_n(\bar{\omega}))_{n=1}^{N(\bar{\omega})}$ is measurable with respect to $\mathcal{F}$. In this way, the only randomness in $\mathrm{E}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right]$ and $\mathrm{Var}\left[\varphi_{n+1} \cdot e_{n+1} \,\middle|\, \mathcal{F}\right]$ comes from the local errors, $(e_n)_{n=1}^{N(\bar{\omega})}$.

Conditional Local Error Representation. In this section, we derive a local error representation that takes into account the fact that the dual is computed backwards and the distribution of the local errors that is relevant to our calculations is therefore not exactly the one given by (3.10), but the distribution given by (3.13).

Consider the sequence $(\bar{X}_n)_{n=0}^{N(\bar{\omega})}$ defined in (3.6). For fixed n, define $\mathcal{F}_n$ as the sigma-algebra

$$\mathcal{F}_n := \sigma\left((\mathcal{Y}_{j,k}(a_j(x_k)\Delta t_k))_{j=1,\dots,J,\,k=1,\dots,n}\right),$$

i.e., the information we obtain by observing the randomness used to generate $\bar{X}_{n+1}$ from x_0. Motivated by dual-weighted expansions (3.5), we want to express the local error representation (3.10) conditional on $\mathcal{F}{:=}\mathcal{F}_{N(\bar{\omega})}$.

At this point, it is convenient to remember a key result for building Poissonian bridges. If X_1 and X_2 are two independent Poisson random variables with parameters λ_1 and λ_2, respectively, we have that $X_1 \,\big|\, X_1 + X_2{=}k$ is a binomial random variable with parameters k and $\lambda_1/(\lambda_1{+}\lambda_2)$.

Applying this observation to the decomposition $\mathcal{Y}_{j,n}(a_j(x_n)\Delta t_n){=}\mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}){+}\mathcal{Q}'_{j,n}(a_j(x_n)\frac{\Delta t_n}{2})$, we conclude that the conditional distribution of $\mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2})$ given $\mathcal{F}_n$, *i.e.*, $\mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) \,\big|\, \mathcal{F}_n$, is binomial with parameters $\mathcal{Y}_{j,n}$ and $1/2$.

Define now the sigma-algebra, $\mathcal{G}_n$, as

$$\mathcal{G}_n{:=}\sigma\left((\mathcal{Q}_{j,n}(a_j(x_n)\frac{\Delta t_n}{2}) \,\big|\, \mathcal{F}_n)_{j=1}^{J}\right).$$

Applying the same argument to $\mathcal{P}''_{j,n}$, defined in (3.9), we conclude that

$$\mathcal{P}''_{j,n} \,\big|\, \{\mathcal{F}_n, \mathcal{G}_n\} \sim \text{binomial}\left(\mathcal{Y}_{j,n}{-}\mathcal{Q}_{j,n}, \frac{c_j}{a_j(x_n)}\right).$$

From the definition of $Z_{n+1}{=}x_n{+}\sum_j \nu_j \mathcal{Q}_{j,n}$ in (3.7), we conclude that

$$\mathcal{R}'_{j,n} \,\big|\, \mathcal{G}_n \sim \text{Poisson}\left((a_j(Z_{n+1}) - m_j)\frac{\Delta t_n}{2}\right).$$

Notice that, by construction, $\mathcal{P}''_{j,n} \,\big|\, \{\mathcal{F}_n, \mathcal{G}_n\}$ and $\mathcal{R}'_{j,n} \,\big|\, \mathcal{G}_n$ are independent random variables. Since $c_j{=}-\Delta a_{j,n}\mathbf{1}_{\{\Delta a_{j,n}<0\}}$ and $a_j(Z_{n+1}){-}m_j{=}\Delta a_{j,n}\mathbf{1}_{\{\Delta a_{j,n}\geq 0\}}$, we can express the conditional local error as

(3.13)

$$e_{n+1} \,\big|\, \{\mathcal{F}_n, \mathcal{G}_n\} = \sum_j \nu_j \left(\mathcal{R}'_{j,n}\left(\Delta a_{j,n}\frac{\Delta t_n}{2}\right)\mathbf{1}_{\{\Delta a_{j,n}\geq 0\}} - \mathcal{P}''_{j,n}\left(\mathcal{Y}_{j,n}{-}\mathcal{Q}_{j,n}, \frac{-\Delta a_{j,n}}{a_j(x_n)}\right)\mathbf{1}_{\{\Delta a_{j,n}<0\}}\right)$$

in the distribution sense. For instance, we can easily compute the expectation of $e_{n+1} \,\big|\, \{\mathcal{F}_n, \mathcal{G}_n\}$ as follows:

$$\mathrm{E}\left[e_{n+1} \,\big|\, \{\mathcal{F}_n, \mathcal{G}_n\}\right] = \sum_j \nu_j \Delta a_{j,n}\left(\frac{\Delta t_n}{2}\mathbf{1}_{\{\Delta a_{j,n}\geq 0\}} + \frac{\mathcal{Y}_{j,n}{-}\mathcal{Q}_{j,n}}{a_j(x_n)}\mathbf{1}_{\{\Delta a_{j,n}<0\}}\right).$$

Taking into account that the joint distribution of $(\mathcal{Q}_{j,n})_{j=1}^{J} \,|\, \mathcal{F}_n$ is given by

$$\mathrm{P}\left((\mathcal{Q}_{j,n}=q_{j,n})_{j=1}^{J} \,|\, \mathcal{F}_n\right) = 2^{-\sum_j \mathcal{Y}_{j,n}} \prod_{j=1}^{J} \frac{\mathcal{Y}_{j,n}!}{q_{j,n}!(\mathcal{Y}_{j,n}-q_{j,n})!}, \quad 0 \le q_{j,n} \le \mathcal{Y}_{j,n},$$

we can exactly compute the expected value and the variance of $v_{n+1}\cdot e_{n+1} \,|\, \mathcal{F}_n$ for any given deterministic vector, v_{n+1}. Notice that given $\mathcal{F}$, the sequence $(\bar{X}_n)_{n=0}^{N(\bar{\omega})}$ is deterministic and, as a consequence, the sequence $(\varphi_n)_{n=1}^{N(\bar{\omega})} \,|\, \mathcal{F}$ is also a deterministic sequence of vectors. We can thus compute

$$\mathrm{E}\left[\sum_n \varphi_{n+1}\cdot e_{n+1} \,|\, \mathcal{F}\right] \text{ and } \mathrm{Var}\left[\sum_n \varphi_{n+1}\cdot e_{n+1} \,|\, \mathcal{F}\right] \tag{3.14}$$

exactly and proceed as stated at the beginning of this section. However, trying to develop computable expressions from (3.13) has two main disadvantages: i) it may lead to computationally demanding procedures, especially for systems with many reaction channels or in regimes with high activity; ii) it may be affected by the variance associated with the randomness in $\mathcal{F}_n$ and $\mathcal{G}_n$.

Deriving a Formula for $\hat{\mathcal{V}}_\ell$. In this section, we derive the formula (3.4). Our goal is to find computable approximations of (3.14), where the underlying sigma-algebra, $\mathcal{F}$, is just the information gathered by observing the coarse path, $\bar{X}$. This means that our formula should not depend explicitly on the knowledge of the random variables that generate $\mathcal{F}_n$ and $\mathcal{G}_n$. At this point, it is important to recall the comments in Section 3.3; that is, the sequence $(\varphi_n(\bar{\omega}))_{n=1}^{N(\bar{\omega})}$ is measurable with respect to $\mathcal{F}$. This implies that, for all n, φ_{n+1} is independent of $\mathcal{G}_n$. Hereafter, for notational convenience, we omit writing explicitly the conditioning on $\mathcal{F}$ in our formulae.

It turns out that the leading order terms of the conditional moments obtained from (3.13) are essentially the same as those computed from (3.10). We will then derive (3.4) from (3.10). Using the notation from Section 3.3, we have that

$$(\varphi_{n+1}\cdot e_{n+1}) = \sum_j f_{j,n}\left(\mathcal{R}'_{j,n}(\Delta a_{j,n}\frac{\Delta t_n}{2})\mathbf{1}_{\{\Delta a_{j,n}>0\}} - \mathcal{P}''_{j,n}(-\Delta a_{j,n}\frac{\Delta t_n}{2})\mathbf{1}_{\{\Delta a_{j,n}<0\}}\right).$$

By the tower property, we obtain

$$\mathrm{E}\left[(\varphi_{n+1}\cdot e_{n+1})\right] = \mathrm{E}\left[\mathrm{E}\left[(\varphi_{n+1}\cdot e_{n+1}) \,|\, \mathcal{G}_n\right]\right] = \frac{\Delta t_n}{2}\sum_j f_{j,n}\mathrm{E}\left[\Delta a_{j,n}\right].$$

Now let us consider the first-order Taylor expansion:

$$\begin{aligned}
\Delta a_{j,n} &:= a_j(x_n + \sum_i \nu_i\, \mathcal{Q}_{i,n}(a_i(x_n)\Delta t_n/2)) - a_j(x_n)\\
&\approx (\nabla a_j(x_n)\cdot \sum_i \nu_i\, \mathcal{Q}_{i,n}(a_i(x_n)\Delta t_n/2))\\
&= \sum_i (\nabla a_j(x_n)\cdot \nu_i)\mathcal{Q}_{i,n}(a_i(x_n)\Delta t_n/2).
\end{aligned}$$

Since $\mathcal{Q}_{i,n}(a_i(x_n)\Delta t_n/2) \sim \mathrm{Poisson}(a_i(x_n)\Delta t_n/2)$, we have that $\mathrm{E}\left[\Delta a_{j,n}\right] = \mu_{j,n}$ and $\mathrm{Var}\left[\Delta a_{j,n}\right] = \sigma^2_{j,n}$. Thus,

$$\mathrm{E}\left[(\varphi_{n+1}\cdot e_{n+1})\right] \approx \frac{\Delta t_n}{2}\sum_j f_{j,n}\,\mu_{j,n}.$$

Now, we use again the tower property for the variance:

$$\mathrm{Var}\left[(\varphi_{n+1}\cdot e_{n+1})\right] = \mathrm{Var}\left[\mathrm{E}\left[(\varphi_{n+1}\cdot e_{n+1})\,\middle|\,\mathcal{G}_n\right]\right] + \mathrm{E}\left[\mathrm{Var}\left[(\varphi_{n+1}\cdot e_{n+1})\,\middle|\,\mathcal{G}_n\right]\right].$$

We then immediately obtain

$$\mathrm{Var}\left[\mathrm{E}\left[(\varphi_{n+1}\cdot e_{n+1})\,\middle|\,\mathcal{G}_n\right]\right] \approx \frac{(\Delta t_n)^3}{8}\sum_{j,j'} f_{j,n}f_{j',n}\sum_i (\nabla a_j(x_n)\cdot\nu_i)(\nabla a_{j'}(x_n)\cdot\nu_i)a_i(x_n),$$

$$\mathrm{E}\left[\mathrm{Var}\left[(\varphi_{n+1}\cdot e_{n+1})\,\middle|\,\mathcal{G}_n\right]\right] \approx \frac{\Delta t_n}{2}\sum_j f_{j,n}^2\,\mathrm{E}\left[\Delta a_{j,n}\,\mathrm{sgn}(\Delta a_{j,n})\right].$$

Let us consider the case where $a_i(x_n)\Delta t_n/2$ is large enough for all i. It is well known that a Poisson random variable, $\mathcal{Q}(\lambda)$, is well approximated by a Gaussian random variable, $N(\lambda,\lambda)$, for moderate values of λ, say $\lambda{>}10$. Since $\mathcal{Q}_{i,n}(a_i(x_n)\Delta t_n/2) \sim \mathrm{Poisson}(a_i(x_n)\Delta t_n/2)$, we have that, when $a_i(x_n)\Delta t_n/2$ is large enough for all i, $\Delta a_{j,n} \approx N(\mu_{j,n},\sigma_{j,n}^2)$. Consider a Gaussian random variable Z with parameters μ and $\sigma^2{>}\,0$. Then,

$$\begin{aligned}\mathrm{E}\left[(\mu+\sigma Z)\mathbf{1}_{\{\mu+\sigma Z>0\}}\right] &= \mu\mathrm{P}\left(\mu+\sigma Z>0\right) + \frac{\sigma}{\sqrt{2\pi}}\int_{-\mu/\sigma}^{+\infty} z\exp\left(-z^2/2\right)dz \\ &= \mu(1-\Phi(-\mu/\sigma)) + \frac{\sigma}{\sqrt{2\pi}}\exp\left(-(\mu/\sigma)^2/2\right).\end{aligned} \tag{3.15}$$

From (3.15), we immediately get

$$\begin{aligned}\mathrm{E}\left[\Delta a_{j,n}\mathbf{1}_{\{\Delta a_{j,n}>0\}}\right] &\approx \mu_{j,n}\left(1-p_{j,n}\right) + \frac{\sigma_{j,n}}{\sqrt{2\pi}}\exp\left(-\frac{q_{j,n}^2}{2}\right), \\ \mathrm{E}\left[\Delta a_{j,n}\mathbf{1}_{\{\Delta a_{j,n}<0\}}\right] &\approx \mu_{j,n}\,p_{j,n} - \frac{\sigma_{j,n}}{\sqrt{2\pi}}\exp\left(-\frac{q_{j,n}^2}{2}\right).\end{aligned} \tag{3.16}$$

By subtracting the expressions in (3.16), we obtain

$$\mathrm{E}\left[\mathrm{Var}\left[(\varphi_{n+1}\cdot e_{n+1})\,\middle|\,\mathcal{G}_n\right]\right] \approx \frac{\Delta t_n}{2}\sum_j f_{j,n}^2\left(\tilde{\mu}_{j,n}+\tilde{\sigma}_{j,n}\right). \tag{3.17}$$

Let us now consider the case where $a_i(x_n)\Delta t_n/2$ is close to zero for some i. We can bound the expression $\mathrm{E}\left[\Delta a_{j,n}\,\mathrm{sgn}(\Delta a_{j,n})\right]$ by $\mathrm{E}\left[|\Delta a_{j,n}|\right]$ and also $\sqrt{\mathrm{E}\left[(\Delta a_{j,n})^2\right]}$. It is easy to see that $\mathrm{E}\left[|\Delta a_{j,n}|\right] \le \bar{\mu}_{j,n}$. Regarding $\mathrm{E}\left[(\Delta a_{j,n})^2\right]$, it can be approximated by

$$\mathrm{E}\left[\sum_{i,i'}(\nabla a_j(x_n)\cdot\nu_i)\left(\nabla a_j(x_n)\cdot\nu_{i'}\right)\mathcal{Q}_i\,\mathcal{Q}_{i'}\right] = \sum_{i,i'}(\nabla a_j(x_n)\cdot\nu_i)\left(\nabla a_j(x_n)\cdot\nu_{i'}\right)\mathrm{E}\left[\mathcal{Q}_i\,\mathcal{Q}_{i'}\right].$$

Since

$$\mathrm{E}\left[\mathcal{Q}_i\,\mathcal{Q}_{i'}\right] = \frac{(\Delta t_n)^2}{4}a_i(x_n)a_{i'}(x_n)\mathbf{1}_{i\neq i'} + \left(a_i(x_n)\frac{\Delta t_n}{2} + \left(a_i(x_n)\frac{\Delta t_n}{2}\right)^2\right)\mathbf{1}_{i=i'}, \tag{3.18}$$

we can rearrange terms and approximate $\mathrm{E}\left[(\Delta a_{j,n})^2\right]$ by $\mu_{j,n}^2 + \sigma_{j,n}^2$.

We conclude that $\mathrm{E}\left[\Delta a_{j,n}\,\mathrm{sgn}(\Delta a_{j,n})\right]$ can be bounded by $m_{j,n}$, which has been defined as $\min\{\bar{\mu}_{j,n}, \sqrt{\mu_{j,n}^2 + \sigma_{j,n}^2}\}$.

REMARK 3.2. *Formula (3.4) can be considered as an initial, relatively successful attempt to estimate $\mathcal{V}_\ell$, but there is still room for improvement. The main problem is the lack of sharp concentration inequalities for linear combinations of independent Poisson random variables. With the numerical examples, we show that the efficiency index of the formula is acceptable for our estimation purposes.*

REMARK 3.3. *We are assuming that only tau-leap steps are taken, but in our hybrid algorithms, some steps can be exact, and, hence, do not contribute to the local error. For that reason, we include the indicator function of the tau-leap step, $\mathbf{1}_{TL}$, in the estimator, $\hat{\mathcal{V}}_\ell$.*

REMARK 3.4. *The dual-weighted residual approach makes the estimation of $\mathcal{V}_\ell$ feasible. In our numerical experiments, we found that, using the same number of simulated coupled hybrid paths, the variance of $\hat{\mathcal{V}}_\ell$ is much smaller than the variance of $\mathrm{Var}\left[g_\ell - g_{\ell-1}\right]$, estimated by a standard Monte Carlo. Note that $\hat{\mathcal{V}}_\ell$ can be computed using only single-level hybrid paths at level $\ell-1$. In the upper right panel of Figure 5.8, we can see that due to the hybrid nature of the simulated paths, it is not possible to predict where the variance of $g_\ell - g_{\ell-1}$ will enter into a superlinear regime. Thus, by extrapolating the $\mathrm{Var}\left[g_\ell - g_{\ell-1}\right]$ from the coarser levels, we may overestimate the values of $\mathrm{Var}\left[g_\ell - g_{\ell-1}\right]$ for the deepest levels.*

4. Estimation Procedure. In this section, we present a procedure that estimates $\mathrm{E}\left[g(X(T))\right]$ within a given prescribed relative tolerance, $TOL>0$, with high probability. The process contains three phases:

Phase I Calibration of virtual machine-dependent quantities.

Phase II Solution of the work optimization problem: we obtain the total number of levels, L, and the sequences $(\delta_\ell)_{\ell=0}^L$ and $(M_\ell)_{\ell=0}^L$, *i.e.*, the one-step exit probability bounds and the required number of simulations at each level. We recall that in Section 1.6, we defined $\Delta t_\ell := \Delta t_0 R^{-\ell}$, where $R > 1$ is a given integer constant. For that reason, to define the whole sequence of meshes, $(\Delta t_\ell)_{\ell=0}^L$, we simply need to define the size of the coarsest mesh, Δt_0.

Phase III Estimation of $\mathrm{E}\left[g(X(T))\right]$.

4.1. Phase I. In this section, we describe the estimation of several constants, C_1, C_2, C_3 and K_1, and functions, C_P and K_2, that allow us to model the expected computational work (or just work), measured in terms of the runtime of hybrid paths, see definitions (4.1) and (4.2). Those quantities are virtual machine dependent; that is, they are dependent on the computer system used for running the simulations and also on the implementation language. Those quantities are also off-line estimated; that is, we need to estimate them only once for each virtual machine on which we want to run the hybrid method.

Constants C_1, C_2, and C_3 reflect the average execution times of each logical path of Algorithm 2. We have that C_1 and C_2 reflect the work associated with the two different types of steps in the MNRM. Constant C_3 reflects the work needed for computing the Chernoff tau-leap size, τ_{Ch}. Finally, when we perform a tau-leap step, we have the work needed for simulating Poisson random variates, which is modeled by the function C_P [19]. This function has two constants that are also virtual machine dependent.

The constant, K_1, and the function, $K_2 \equiv K_2(x,\delta)$, defined through C_1, C_2, and C_3, were introduced in Section 1.5.

4.2. Phase II. In this section, we set and solve the work optimization problem. Our objective function is the expected total work of the MLMC estimator, $\mathcal{M}_L$, defined in (3.1), *i.e.*,

$$\sum_{\ell=0}^{L} \psi_\ell M_\ell,$$

where L is the maximum level (deepest level), ψ_0 is the expected work of a single-level path at level 0, and ψ_ℓ, for $\ell \geq 1$, is the expected computational work of two coupled paths at levels $\ell-1$ and ℓ. Finally, M_0 is the number of single-level paths at level 0, and M_ℓ, for $\ell \geq 1$, is the number of coupled paths at levels $\ell-1$ and ℓ.

Let us now describe in detail the quantities, $(\psi_\ell)_{\ell=0}^{L}$. For ℓ=0, Algorithm 12 generates a single hybrid path. The building block of a single hybrid path is Algorithm 2, which adaptively determines whether to use an MNRM step or a tau-leap one. According to this algorithm, there are two ways of taking an MNRM step, depending on the logical conditions, $K_1/a_0(x)>T_0-t$ and $K_2/a_0(x)>\tau_{Ch}$. Given one particular hybrid path, let $N_{K1}(\Delta t_0, \delta_0)$ be the number of MNRM steps such that $K_1/a_0(x)>T_0-t$ is true, and let $N_{K2}(\Delta t_0, \delta_0)$ be the number of MNRM steps such that $K_1/a_0(x)>T_0-t$ is false and $K_2/a_0(x)>\tau_{Ch}$ is true. When a Chernoff tau-leap step is taken, we have constant work, C_3, and variable work computed with the aid of C_P. Then, the expected work of a single hybrid path, at level $\ell = 0$, is

$$\begin{aligned}(4.1)\qquad \psi_0 := {} & C_1 \mathrm{E}\left[N_{K1}(\Delta t_0, \delta_0)\right] + C_2 \mathrm{E}\left[N_{K2}(\Delta t_0, \delta_0)\right] + C_3 \mathrm{E}\left[N_{\mathrm{TL}}(\Delta t_0, \delta_0)\right] \\ & + \sum_{j=1}^{J} \mathrm{E}\left[\int_{[0,T]} C_P(a_j(\bar{X}_0(s))\tau_{Ch}(\bar{X}_0(s), \delta_0))\mathbf{1}_{TL}(\bar{X}_0(s))ds\right],\end{aligned}$$

where Δt_0 is the size of the time mesh at level 0 and δ_0 is the exit probability bound at level 0. Therefore, the expected work at level 0 is $\psi_0 M_0$, where M_0 is the total number of single hybrid paths.

For $\ell \geq 1$, we use Algorithm 3 to generate M_ℓ-coupled paths that couple the $\ell-1$ and ℓ levels. Given two coupled paths, let $N_{K1}(\Delta t_{\ell-1}, \delta_{\ell-1})$ and $N_{K1}(\Delta t_\ell, \delta_\ell)$ be the number of exact steps for level $\ell-1$ (coarse mesh) and ℓ (fine mesh), respectively, with associated work C_1. We define $N_{K2}(\Delta t_{\ell-1}, \delta_{\ell-1})$ and $N_{K2}(\Delta t_\ell, \delta_\ell)$ analogously. Then, the expected work of a pair of coupled hybrid paths at levels ℓ and $\ell-1$ is

$$\begin{aligned}(4.2)\qquad \psi_\ell := {} & C_1 \mathrm{E}\left[N_{K1}^{(c)}(\ell)\right] + C_2 \mathrm{E}\left[N_{K2}^{(c)}(\ell)\right] + C_3 \mathrm{E}\left[N_{\mathrm{TL}}^{(c)}(\ell)\right] \\ & + \sum_{j=1}^{J} \mathrm{E}\left[\int_{[0,T]} C_P(a_j(\bar{X}_\ell(s))\tau_{Ch}(\bar{X}_\ell(s), \delta_\ell))\mathbf{1}_{TL}(\bar{X}_\ell(s))ds\right] \\ & + \sum_{j=1}^{J} \mathrm{E}\left[\int_{[0,T]} C_P(a_j(\bar{X}_{\ell-1}(s))\tau_{Ch}(\bar{X}_{\ell-1}(s), \delta_{\ell-1}))\mathbf{1}_{TL}(\bar{X}_{\ell-1}(s))ds\right],\end{aligned}$$

where

$$\begin{aligned} N_{K1}^{(c)}(\ell) &:= N_{K1}(\Delta t_\ell, \delta_\ell) + N_{K1}(\Delta t_{\ell-1}, \delta_{\ell-1}) \\ N_{K2}^{(c)}(\ell) &:= N_{K2}(\Delta t_\ell, \delta_\ell) + N_{K2}(\Delta t_{\ell-1}, \delta_{\ell-1}) \\ N_{\mathrm{TL}}^{(c)}(\ell) &:= N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell) + N_{\mathrm{TL}}(\Delta t_{\ell-1}, \delta_{\ell-1}). \end{aligned}$$

Now, recalling the definitions of the error decomposition, given at the beginning of Section 3.2, we have all the elements to formulate the work optimization problem. Given a relative tolerance, $TOL>0$, we solve

$$\begin{cases} \min_{\{\Delta t_0, L, (M_\ell, \delta_\ell)_{\ell=0}^L\}} \sum_{\ell=0}^L \psi_\ell M_\ell \\ \text{s.t.} \\ \mathcal{E}_{E,L} + \mathcal{E}_{I,L} + \mathcal{E}_{S,L} \leq TOL. \end{cases} \tag{4.3}$$

It is natural to consider the following family of auxiliary problems indexed on $L\geq 1$, where we assume for now that the double sequence, $(\Delta t_\ell, \delta_\ell)_{\ell=0}^L$, is known:

$$\begin{cases} \min_{(M_\ell \geq 1)_{\ell=0}^L} \sum_{\ell=0}^L \psi_\ell M_\ell \\ \text{s.t.} \\ \mathcal{E}_{I,L} + C_A \sqrt{\sum_{\ell=0}^L \frac{\mathcal{V}_\ell}{M_\ell}} \leq TOL - TOL^2, \end{cases} \tag{4.4}$$

where we have $C_A \geq 2$ to guarantee an asymptotic confidence level of at least 95%.

Let us assume for now that we know ψ_ℓ, $\mathcal{V}_\ell$ and $\mathcal{E}_{I,\ell}$, for $\ell = 0, 1, \ldots, L$. Let L_0 be the smallest value of L such that $\mathcal{E}_{I,L} < TOL - TOL^2$. This value exists and it is finite since the discretization error, $\mathcal{E}_{I,L}$, tends to zero as L goes to infinity. For each $L \geq L_0$, define $w_L := \sum_{\ell=0}^L \psi_\ell M_\ell^*$, where the sequence $(M_\ell^*)_{\ell=0}^L$ is the solution of the problem (4.4). It is worth mentioning that $(M_\ell^*)_{\ell=0}^L$ is quickly obtained as the solution of the following Karush-Kuhn-Tucker problem (see, *e.g.*, [18]):

$$\begin{cases} \min_{(M_\ell \geq 1)_{\ell=0}^L} \sum_{\ell=0}^L \psi_\ell M_\ell \\ \text{s.t.} \\ \sum_{\ell=0}^L \frac{\mathcal{V}_\ell}{M_\ell} \leq R \end{cases} . \tag{4.5}$$

We do not develop here all the calculations, but a pseudo code is given in Algorithm 11.

Let us now analyze two extreme cases: i) for L such that $\mathcal{E}_{I,L}$ is less but very close to $TOL - TOL^2$, we have that $\sum_{\ell=0}^L \mathcal{V}_\ell / M_\ell^*$ is a very small number. As a consequence, we obtain large values of M_ℓ^* and, hence, a large value of w_L. By adding one more level, *i.e.*, $L \leftarrow L+1$, we expect a larger gap between $\mathcal{E}_{I,L}$ and TOL_0; that means that we expect a larger value of $\sum_{\ell=0}^L \mathcal{V}_\ell / M_\ell^*$ that may lead to smaller values of M_ℓ^*. We observe that, in spite of adding one more term to w_L, this leads to a smaller value of w_L. ii) At the other extreme, a large value of L is associated with large values of ψ_L and therefore with large values of w_L.

This informal 'extreme case analysis' has been confirmed by our numerical experiments (see, for instance, Figures 5.2 and 5.8 (lower-right)), which allow us to conjecture that the sequence $(w_L)_{L=L_0}^{+\infty}$ is a convex function of L and, hence, that it has a unique optimal value achieved at a certain L^*. A pseudo algorithm to find L^* could be to start computing w_{L_0} and w_{L_0+1}. If $w_{L_0+1} \geq w_{L_0}$, we accept $L^* = L_0$; otherwise, we proceed to computing the next term of the sequence, $(w_L)_{L=L_0}^{+\infty}$. If, for some p, we have $w_{L_{p+1}} \geq w_{L_p}$, we accept $L^* = L_p$. Of course, we can stop even if $w_{L_{p+1}} < w_{L_p}$, but the difference $\left| w_{L_{p+1}} - w_{L_p} \right|$ is sufficiently small. In this last case, we accept $L^* = L_{p+1}$.

4.2.1. Computational Complexity. At this point, we have all the necessary elements to establish a key point of this work, the computational complexity of the multilevel hybrid Chernoff tau-leap method.

Let us now analyze the optimal amount of work at level L, w_L, as a function of the given relative tolerance, TOL. For simplicity, let us assume that $M_\ell^*>1$, $\ell=0,...,L$. In this case, the optimal number of samples at level ℓ is given by

$$M_\ell^*=(C_A/\theta)^2 TOL^{-2}\sqrt{\mathcal{V}_\ell/\psi_\ell}\sum_{\ell=0}^{L}\sqrt{\mathcal{V}_\ell\psi_\ell},$$

for some $\theta\in(0,1)$. In fact, θ is the proportion of the tolerance, TOL, that our computational cost optimization algorithm selects for the statistical error, $\mathcal{E}_{S,L}$. In our algorithms, we impose $\theta\geq 0.5$; however, our numerical experiments always select a larger value (see Figures 5.3 and 5.9).

By substituting M_ℓ^* into the total work formula, w_L, we conclude that the optimal expected work, conditional on θ, is given by

$$\mathrm{E}\left[w_L^*(TOL)\,\middle|\,\theta\right]=\left(\frac{C_A}{\theta}\sum_{\ell=0}^{L(\theta)}\sqrt{\mathcal{V}_\ell\psi_\ell}\right)^2 TOL^{-2}.$$

Due to the constraint, $\theta\geq 0.5$, we have that

$$w_L^*(TOL)\leq \sup_L\left\{\left(2\,C_A\sum_{\ell=0}^{L}\sqrt{\mathcal{V}_\ell\psi_\ell}\right)^2\right\}TOL^{-2}.$$

Let us consider the series $\sum_{\ell=0}^{\infty}\sqrt{\mathcal{V}_\ell\psi_\ell}$. First, observe that the expected computational work per path at level ℓ, ψ_ℓ, is bounded by a multiple of the expected computational work of the MNRM (see Section 1.2), *i.e.*, $K\psi_{\mathrm{MNRM}}$. In our numerical experiments, we observe that taking K around 3 is enough. Therefore, $\sum_{\ell=0}^{\infty}\sqrt{\mathcal{V}_\ell\psi_\ell}\leq\sqrt{K\psi_{\mathrm{MNRM}}}\sum_{\ell=0}^{\infty}\sqrt{\mathcal{V}_\ell}$. Observe that, by construction, $\mathcal{V}_\ell\to 0$, superlinearly. More specifically, it satisfies the bound $\mathcal{V}_\ell=\mathcal{O}\left(\Delta t_\ell\right)\leq C\Delta t_0(1/2)^\ell$ for some positive constant C. Therefore, the series $\sum_{\ell=0}^{\infty}\sqrt{\mathcal{V}_\ell}$ is dominated by the geometric series $\sum_{\ell=0}^{\infty}(1/\sqrt{2})^\ell<\infty$. We conclude that $\sup_L\{\sum_{\ell=0}^{L}\sqrt{\mathcal{V}_\ell\psi_\ell}\}$ is bounded and, therefore, the expected computational complexity of the multilevel hybrid Chernoff tau-leap method is $w_L^*(TOL)=\mathcal{O}\left(TOL^{-2}\right)$.

4.2.2. Some Comments on the Algorithms for Phase II. In Algorithm 7, we propose an iterative method to obtain an approximate solution to the problem (4.3). Notice that we are assuming that there are at least two levels in the multilevel hierarchy, *i.e.*, $L\geq 1$.

To solve the problem (4.3), we bound the global exit error, $\mathcal{E}_{E,L}$, by TOL^2. More specifically, we choose δ_L to be sufficiently small such that

$$|\mathcal{A}\left(g_L;\cdot\right)|\;\delta_L\;\mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_L,\delta_L);\cdot\right)<TOL^2. \tag{4.6}$$

At this point, it is crucial to observe that if we impose the condition (4.6) on any level $\ell<L$, then we are unnecessarily enforcing a dependence of δ_ℓ on TOL. This dependence may result in very small values of δ_ℓ, which in turn may increase the expected number of exact steps and tau-leap steps at level ℓ, implying a larger expected computational work at level ℓ. In the appendix of [19], we proved that, when δ_ℓ tends to zero, the expected values of the number of tau-leap steps at level ℓ go to zero, and therefore our hybrid MLMC strategy would converge to the SSA method without

the desired reduction in computational work. To avoid the dependence of $(\delta_\ell)_{\ell=0}^{L-1}$ on TOL, we adopt a different strategy based on the following decomposition:

$$\begin{aligned}\mathcal{V}_\ell = \mathrm{Var}\left[g_\ell \mathbf{1}_{A_\ell} - g_{\ell-1}\mathbf{1}_{A_{\ell-1}}\right] &= \mathrm{Var}\left[g_\ell - g_{\ell-1} \,\middle|\, A_\ell \cap A_{\ell-1}\right] \mathrm{P}\left(A_\ell \cap A_{\ell-1}\right) \\ &+ \mathrm{Var}\left[g_\ell \,\middle|\, A_\ell \cap A_{\ell-1}^c\right] \mathrm{P}\left(A_\ell \cap A_{\ell-1}^c\right) \\ &+ \mathrm{Var}\left[g_{\ell-1} \,\middle|\, A_\ell^c \cap A_{\ell-1}\right] \mathrm{P}\left(A_\ell^c \cap A_{\ell-1}\right).\end{aligned}$$

We impose that the first term of the right-hand side dominates the other two. This is because the conditional variances appearing in the last two terms are of order $\mathcal{O}(1)$, while the conditional variance appearing in the first term is of order $\mathcal{O}(\Delta t_\ell)$, and we make our computations with approximations of $\mathcal{V}_\ell$ assuming that $\mathrm{P}\left(A_\ell \cap A_{\ell-1}\right)$ is close to one. We proceed as follows: first, we approximate $\mathrm{P}\left(A_\ell \cap A_{\ell-1}\right)$ by $\mathrm{P}\left(A_\ell\right)\mathrm{P}\left(A_{\ell-1}\right)$; then, we consider $1-\delta_\ell \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell);\cdot\right)$ as an approximate upper bound for $\mathrm{P}\left(A_\ell\right)$ when $\delta_\ell \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell);\cdot\right) \ll 1$. Those considerations lead us to impose

(4.7)
$$\begin{aligned}&\mathrm{Var}\left[g_\ell - g_{\ell-1} \,\middle|\, A_\ell \cap A_{\ell-1}\right] \left(1-\delta_\ell \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell);\cdot\right)\right)\left(1-\delta_{\ell-1} \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_{\ell-1}, \delta_{\ell-1});\cdot\right)\right) > \\ &\mathrm{Var}\left[g_\ell \,\middle|\, A_\ell \cap A_{\ell-1}^c\right] \delta_{\ell-1} \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_{\ell-1}, \delta_{\ell-1});\cdot\right) + \mathrm{Var}\left[g_{\ell-1} \,\middle|\, A_\ell^c \cap A_{\ell-1}\right] \delta_\ell \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell);\cdot\right).\end{aligned}$$

To avoid simultaneous refinements on δ_ℓ and $\delta_{\ell-1}$, based on (4.7), we impose on δ_ℓ the following condition:

$$\hat{\mathcal{V}}_\ell \left(1-\delta_\ell \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell);\cdot\right)\right)^2 > 2\, \mathcal{S}^2\left(g;\cdot\right) \delta_\ell \mathcal{A}\left(N_{\mathrm{TL}}(\Delta t_\ell, \delta_\ell);\cdot\right). \tag{4.8}$$

Algorithms 12 and 7 provide $\mathcal{A}\left(g_\ell;\cdot\right)$, $\mathcal{A}\left(N_{\mathrm{TL}};\cdot\right)$ and the other required quantities. Condition (4.8) does not affect the telescoping sum property of our multilevel estimator, $\mathcal{M}_L$, defined in (3.1), since each level, ℓ, has its own δ_ℓ.

Remark 4.1 (Multilevel estimators used in Algorithm 7). *Although in algorithm 7 we show that the estimations of* $\mathrm{E}\left[g(X(T))\right]$ *and* $\mathrm{Var}\left[g(X(T))\right]$ *are computed using the information from the last level only, in fact we are computing them using a multilevel estimator. We omit the details in the algorithm for the sake of simplicity. For the case of* $\mathrm{E}\left[g(\bar{X}(T))\right]$*, we use the standard mutilevel estimator, and, for the case of* $\mathrm{Var}\left[g(X(T))\right]$*, we use the following telescopic decomposition:*

$$\mathrm{Var}\left[g(\bar{X}_l(T))\right] = \mathrm{Var}\left[g(\bar{X}_0(T))\right] + \sum_{\ell=1}^{l} (\mathrm{Var}\left[g(\bar{X}_\ell(T))\right] - \mathrm{Var}\left[g(\bar{X}_{\ell-1}(T))\right]),$$

where $l > 1$ *is a fixed level. Using the usual variance estimators for each level, we obtain an unbiased multilevel estimator of the variance of* $g(\bar{X})$*. We refer to [5] for details.*

Remark 4.2 (Coupled paths exiting the lattice, $\mathbb{Z}_+^d$). *Algorithm 3 could compute four types of paths. It could happen that no approximate process (the coarse one,* $X_{\ell-1}$*, or the fine one,* X_ℓ*) exits the lattice, which is the most common case. It could also happen that one of the approximate processes exits the lattice. And finally, both approximate processes could exit the lattice. The first case is the most common one and no further explanation is required. We now explain the case when one of processes exits the lattice. Suppose that the coarse one exits the lattice. In that case, until the fine process reaches time* T *or exits the lattice, we still simulate the coupled process by simulating only the fine path using the single-level hybrid algorithm presented in [19].*

If the fine path reaches T, we have that $\mathbf{1}_{A_{\ell-1}} = 0$, and $\mathbf{1}_{A_\ell} = 1$. Vice versa, if the fine process exits and the coarse one reaches T, we have $\mathbf{1}_{A_{\ell-1}} = 1$ and $\mathbf{1}_{A_\ell} = 0$.

REMARK 4.3 (Coupling with an exact path). *Algorithm 7 uses a computational-cost-based stopping criterion. That is, the algorithm stops refining the time mesh when the estimated total computational cost of the multilevel estimator, $\hat{W}_{ML} := \sum_{\ell=0}^{l} \hat{\psi}_\ell M_\ell$, at level l, is greater than the corresponding computational cost for level $l-1$, and only when condition $\hat{\mathcal{E}}_I < TOL - TOL^2$ is already satisfied. In that case, $L^* = l-1$. The latter condition is required for obtaining a solution of the optimization problem (4.5). In our numerical experiments, we observed that the computational cost of two hybrid coupled paths, ψ_ℓ, may be greater than the computational cost of "hybrid-exact" coupled paths; that is, the computational cost of a hybrid path at level $l-1$ coupled with an exact path at level l. That kind of path, used only at the last level, leads to the following unbiased multilevel estimator:*

$$\tilde{\mathcal{M}}_L := \frac{1}{M_0}\sum_{m=1}^{M_0} g_0 \mathbf{1}_{A_0}(\omega_{m,0}) + \sum_{\ell=1}^{L-1}\frac{1}{M_\ell}\sum_{m=1}^{M_\ell}[g_\ell \mathbf{1}_{A_\ell} - g_{\ell-1}\mathbf{1}_{A_{\ell-1}}](\omega_{m,\ell}) + \frac{1}{M_L}\sum_{m=1}^{M_L}[g(X(T)) - g_{\ell-1}\mathbf{1}_{A_{L-1}}](\omega_{m,L}).$$

Therefore, it is possible to add another stopping criterion to Algorithm 7 related to the comparison between the estimated computational cost of two hybrid coupled paths and the computational cost of hybrid-exact coupled paths. Please note that the condition $\delta_L \mathcal{A}(N_{TL,L};\cdot)\,\mathcal{A}(g_L;\cdot) \leq TOL^2$ trivially holds because $\mathcal{A}(N_{TL,L};\cdot)$ is zero in such a case. In our numerical examples, there are no significant computational gains in the estimation phase from using that stopping rule and its corresponding estimator. This alternative hybrid unbiased estimator is inspired by the work of Anderson and Higham [1].

4.3. Phase III. From Phase II, we found that, to compute our multilevel Monte Carlo estimator, $\mathcal{M}_L$, for a given tolerance, we have to run M_0^* single hybrid paths with parameters $(\Delta t_0, \delta_0)$ and M_ℓ^* coupled hybrid paths with parameters $(\Delta t_{\ell-1}, \delta_{\ell-1})$ and $(\Delta t_\ell, \delta_\ell)$, for $\ell = 1, 2, \ldots, L^*$. But, we will follow a slightly different strategy: we run half of the required simulations and use them to update our estimations of the sequences $(\mathcal{E}_{I,\ell})_{\ell=0}^{L^*}$, $(\mathcal{V}_\ell)_{\ell=0}^{L^*}$, and $(\psi_\ell)_{\ell=0}^{L^*}$. Then, we solve the problem (4.4) again and re-calculate the values of M_ℓ^* for all ℓ. We proceed iteratively until convergence. In this way, we take advantage of the information generated by new simulated paths and update the estimations of the sequences of weak errors, computational costs, and variances, obtaining more control over the total work of the method.

5. Numerical Examples. In this section, we present two examples to illustrate the performance of our proposed method, and we compare the results with the single-level approach given in [19]. For bench-marking purposes, we use Gillespie's Stochastic Simulation Algorithm (SSA) instead of the Modified Next Reaction Method (MNRM), because the former is widely used in the literature.

5.1. A Simple Decay Model. The classical radioactive decay model provides a simple and important example for the application of our method. This model has only one species and one first-order reaction,

$$X \xrightarrow{c} \emptyset. \tag{5.1}$$

Its stoichiometric matrix, $\nu \in \mathbb{R}$, and the propensity function, $a : \mathbb{Z}_+ \to \mathbb{R}$, are given by

$$\nu = -1 \text{ and } a(X) = c\,X.$$

Here, we choose $c = 1$, and define $g(x) = x$ as the scalar observable. In this particularly simple example, we have that $\mathrm{E}\left[g(X(T))|X(t) = X_0\right] = X_0 \exp(-c(T-t))$. Consider the initial condition $X_0{=}10^5$ and the final time $T{=}0.5$. In this case, the process starts relatively far from the boundary, *i.e.*, it is a tau-leap dominated setting.

We now analyze an ensemble of five independent runs of the calibration algorithm (Algorithm 7), using different relative tolerances. In Figure 5.1, we show, in the left panel, the total predicted work (runtime) for the single-level hybrid method, for the multilevel hybrid method and for the SSA method, versus the estimated error bound. The multilevel method is preferred over the SSA and the single-level hybrid method for all the tolerances. We also show the estimated asymptotic work of the multilevel method. In the right panel, we show, for different tolerances, the actual work (runtime), using a 20 core Intel GLNXA64 architecture and MATLAB version R2014a.

In Table 5.1, we summarize an ensemble run of the calibration algorithm, where W_{ML} is the average actual computational work of the multilevel estimator (the sum of all the seconds taken to compute the estimation) and W_{SSA} is the corresponding average actual work of the SSA. We compare those values with the corresponding estimations, $\hat{W}_{\mathrm{ML}}$ and $\hat{W}_{\mathrm{SSA}}$.

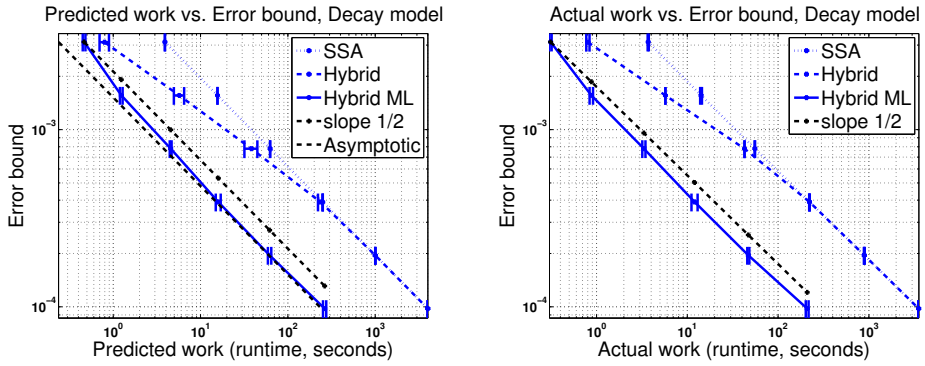

FIG. 5.1. *Left: Predicted work (runtime) versus the estimated error bound for the simple decay model (5.1), with 95% confidence intervals. The multilevel hybrid method is preferred over the SSA and the single-level method for all the tolerances. Right: Actual computational work (runtime) versus the estimated error bound. Notice that the computational complexity has order* $\mathcal{O}\left(TOL^{-2}\right)$.

In Figure 5.2, we can observe how the estimated weak error, $\hat{\mathcal{E}}_{I,\ell}$, and the estimated variance of the difference of the functional between two consecutive levels, $\hat{\mathcal{V}}_\ell$, decrease linearly as we refine the time mesh. This corresponds to the pure tau-leap case since the process, X, remains far from the boundary in $[0, T]$. As expected, the linear relationship for the variance starts at level 1. The estimated total path work, $\hat{\psi}_\ell$, increases as we refine the mesh. Observe that it increases more slowly than linearly. This is because the work needed for generating Poisson random variables becomes less as we refine the time mesh. In the lower right panel, we show the total computational work, only in the cases in which $\hat{\mathcal{E}}_{I,\ell} < TOL{-}TOL^2$.

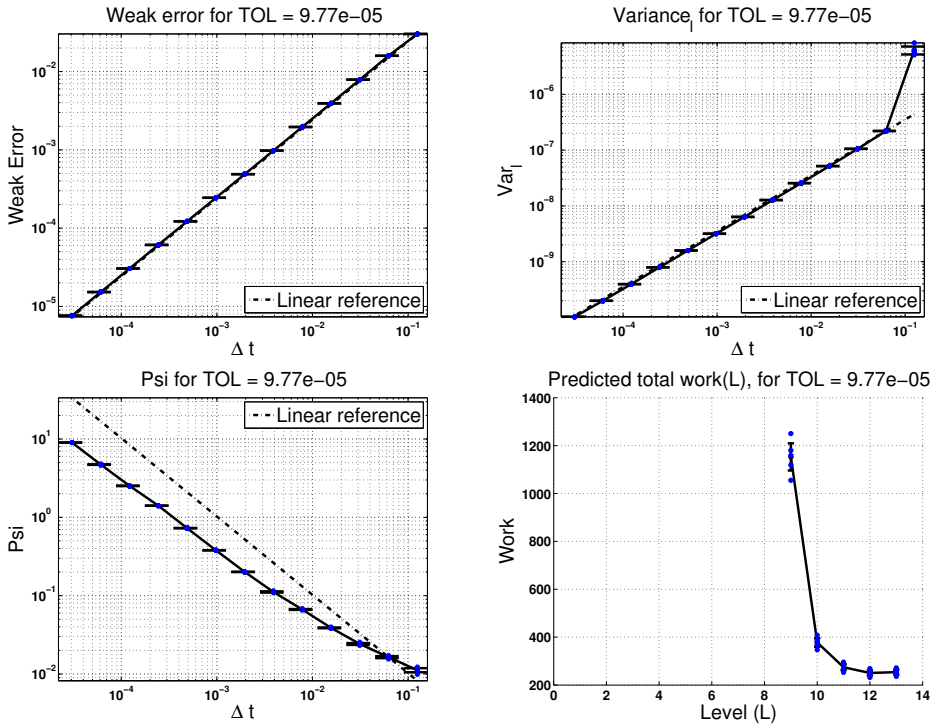

FIG. 5.2. *Upper left: estimated weak error, $\hat{\mathcal{E}}_{I,\ell}$, as a function of the time mesh size, Δt, for the simple decay model (5.1). Upper right: estimated variance of the difference between two consecutive levels, $\hat{\mathcal{V}}_\ell$, as a function of Δt. Lower left: estimated path work, $\hat{\psi}_\ell$, as a function of Δt. Lower right: estimated total computational work, $\sum_{l=0}^{L} \hat{\psi}_l M_l$, as a function of the level, L.*

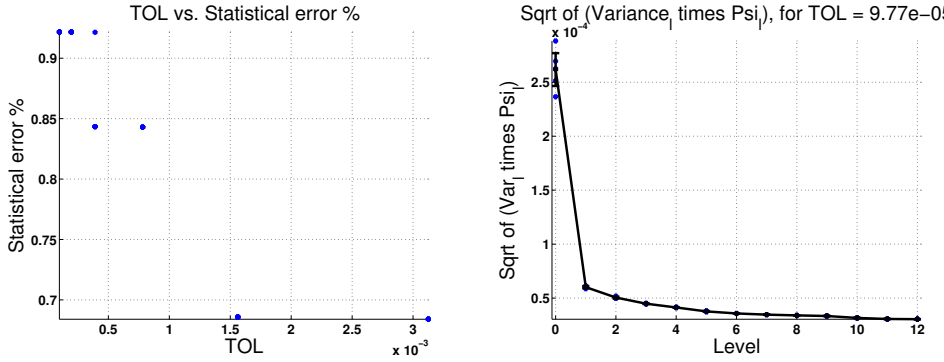

FIG. 5.3. *Left: Percentage of the statistical error over the computational global error, for the simple decay model (5.1). As mentioned in Section 4, it is well above 0.5 for all the tolerances. Right: $\sqrt{\hat{\mathcal{V}}_\ell \hat{\psi}_\ell}$ as a function of ℓ, for the smallest tolerance, which decreases as the level increases. Observe that the contribution of level 0 is less than 50% of the sum of the other levels.*

In Figure 5.4, we show the main outputs of Algorithm 7, δ_ℓ and M_ℓ for $\ell = 0, ..., L^*$, for the smallest considered tolerance. In this case, L^* is 12. We observe that the number of realizations decreases slower than linearly, from levels 1 to L^*-1, until it reaches $M_{L^*}=1$.

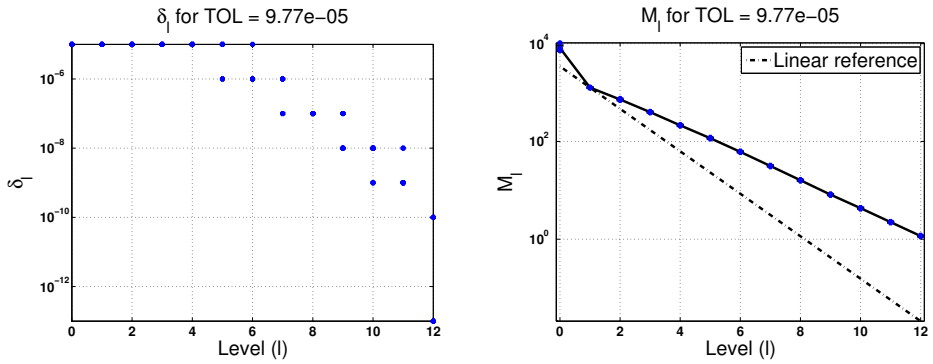

FIG. 5.4. *One-step exit probability bound, δ_ℓ, and M_ℓ for ℓ=0, 1, ..., L^*, for the smallest tolerance, for the simple decay model (5.1).*

TOL	L^*	Min	Max	$\frac{\hat{W}_{ML}}{\hat{W}_{\text{SSA}}}$	Min	Max	$\frac{W_{ML}}{W_{\text{SSA}}}$	Min	Max
3.13e-03	5	5	5	0.03	0.02	0.04	0.03	0.02	0.05
1.56e-03	6	6	6	0.04	0.02	0.10	0.04	0.02	0.13
7.81e-04	8	8	8	0.03	0.02	0.05	0.03	0.02	0.06
3.91e-04	9.2	9	10	0.02	0.02	0.03	0.02	0.01	0.03
1.95e-04	11	11	11	0.02	0.02	0.03	0.02	0.02	0.04
9.77e-05	12	12	12	0.03	0.02	0.03	0.03	0.02	0.03

TABLE 5.1

Details of the ensemble run of Algorithm 7 for the simple decay model (5.1). As an example, the second row of the table indicates that, for a tolerance TOL=$1.56 \cdot 10^{-3}$, six levels are needed. The predicted work of the multilevel hybrid method is, on average, 4% of the predicted work of the SSA method, which coincides with the actual work. Observed minimum and maximum values in the ensemble are also provided.

In the left panel of Figure 5.5, we show the performance of formula (3.4), implemented in Algorithm 10, used to estimate the strong error, $\mathcal{V}_\ell$, defined in Section 3.2. The quotient of $\hat{\mathcal{V}}_\ell$ over a standard Monte Carlo estimate of $\mathcal{V}_\ell$ is almost 1 for the first ten levels. At levels 11 and 12, we obtain 0.99 and 0.91, respectively. Both quantities are estimated using a coefficient of variation less than 5%, but there is a remarkable difference in terms of computational work in favor of our dual-weighted estimator. In the right panel of the same figure, we show the estimated variance of $\mathcal{V}_\ell$, computed by dual-weighted estimation (3.4) and computed by direct sampling. Observe that, in this case, the computational savings may be up to order $\mathcal{O}\left(10^5\right)$.

In the simulations, we observed that, as we refine TOL, the optimal number of levels approximately increases logarithmically, which is a desirable feature. We fit the model $L^* = a + b\log(TOL^{-1})$, obtaining b=−2.11 and a=−7.3.

The QQ-plot in Figure 5.6 shows, for the smallest considered TOL, 10^3 independent realizations of the multilevel estimator, $\mathcal{M}_L$ (defined by (3.1)). Those 10^3 points are generated using 5 sets of parameters given by an independent run of the calibration algorithm (Algorithm 7). This plot, complemented with a Shapiro-Wilk normality test, validates our assumption about the Gaussian distribution of the statistical error. Observe that the estimates are concentrated around the theoretical

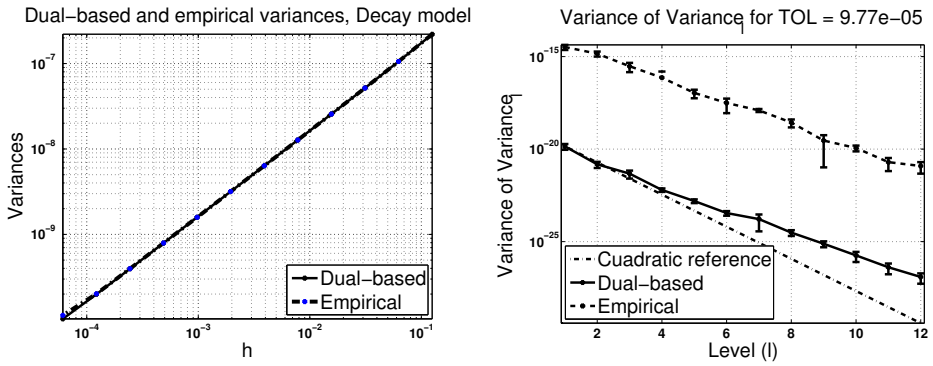

FIG. 5.5. *Left: performance of the formula* (3.4) *as a strong error estimate, for the simple decay model* (5.1)*. Here,* $h{=}\Delta t$*. Right: estimated variance of* $\mathcal{V}_\ell$ *with 95% confidence intervals.*

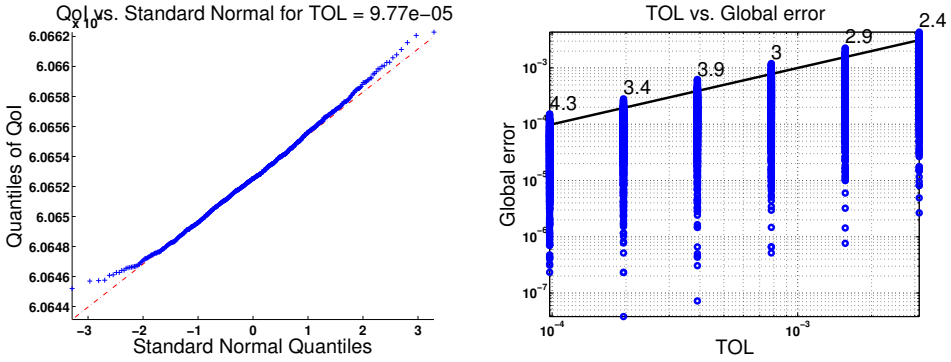

FIG. 5.6. *Left: QQ-plot for the hybrid Chernoff MLMC estimates,* $\mathcal{M}_L$*, in the simple decay model* (5.1)*. Also, we performed a Shapiro-Wilk normality test, and we obtained a p-value of* 0.0105*. Right: TOL versus the actual computational error. The numbers above the straight line show the percentage of runs that had errors larger than the required tolerance. We observe that in all cases, except for the smallest tolerance, the computational error follows the imposed tolerance with the expected confidence of 95%.*

value $X_0 \exp(-c(T-t)) = 10^5 \exp(-0.5) \approx 6.0653e+04$. In the same figure, we also show TOL versus the actual computational error. It can be seen that the prescribed tolerance is achieved with the required confidence of 95%, in all the tolerances.

5.2. Gene Transcription and Translation [1]. This model has five reactions,

$$\begin{aligned} \emptyset \xrightarrow{c_1} R, \quad R \xrightarrow{c_2} R + P \\ 2P \xrightarrow{c_3} D, \quad R \xrightarrow{c_4} \emptyset \\ P \xrightarrow{c_5} \emptyset \end{aligned} \tag{5.2}$$

described respectively by the stoichiometric matrix and the propensity function

$$\nu = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -2 & 1 \\ -1 & 0 & 0 \\ 0 & -1 & 0 \end{pmatrix} \text{ and } a(X) = \begin{pmatrix} c_1 \\ c_2 R \\ c_3 P(P-1) \\ c_4 R \\ c_5 P \end{pmatrix},$$

where $X(t) = (R(t), P(t), D(t))$, and c_1=25, c_2=10^3, c_3=0.001, c_4=0.1, and c_5=1. In the simulations, the initial condition is $(0, 0, 0)$ and the final time is T=1. The observable is given by $g(X) = D$. We observe that the abundance of the mRNA species, represented by R, is close to zero for $t \in [0, T]$. However, as we point out in [19], the reduced abundance of one of the species is not enough to ensure that the SSA method should be used.

We now analyze an ensemble of five independent runs of the calibration algorithm (Algorithm 7), using different relative tolerances. In Figure 5.7, we show, in the left panel, the total predicted work (runtime) for the single-level hybrid method, for the multilevel hybrid method and for the SSA method, versus the estimated error bound. We also show the estimated asymptotic work of the multilevel method. Again, the multilevel hybrid method outperforms the others and we remark that the observed computational work of the multilevel method is of order $\mathcal{O}\left(TOL^{-2}\right)$.

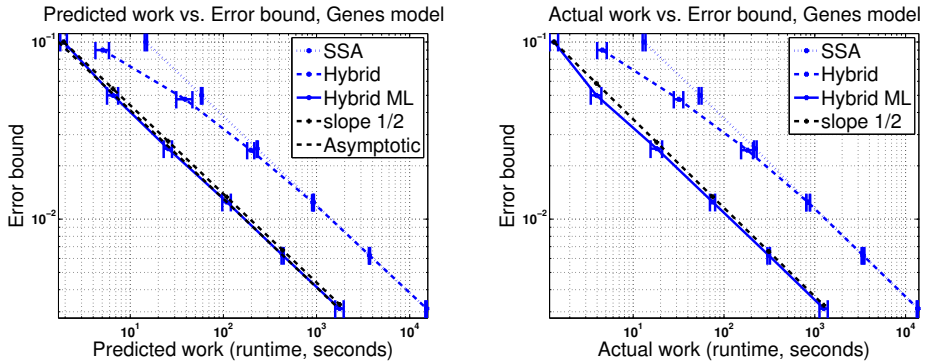

FIG. 5.7. *Left: Predicted work (runtime) versus the estimated error bound for the gene transcription and translation model (5.2). The hybrid method is preferred over the SSA for the first three tolerances only. The multilevel hybrid method is preferred over the SSA and the single-level method for all the tolerances. Right: Actual work (runtime) versus the estimated error bound.*

In Figure 5.8, we can observe how the estimated weak error decreases linearly for the coarser time meshes, but, as we continue refining the time mesh, it quickly decreases towards zero. In the case of the estimated variance, $\hat{\mathcal{V}}_\ell$, it decreases faster than linearly, and it also quickly decreases towards zero afterwards. This is a consequence of the transition from a hybrid regime to a pure exact one. The estimated total path work, $\hat{\psi}_\ell$, increases sublinearly as we refine the mesh. Note that $\hat{\psi}_\ell$ reaches a maximum, which corresponds to a SSA-dominant regime. In the lower right panel, we show the total computational work only in the cases in which $\hat{\mathcal{E}}_{I,\ell} < TOL - TOL^2$.

In Figure 5.10, we show the main outputs of Algorithm 7, δ_ℓ and M_ℓ for $\ell = 0, ..., L^*$, for the smallest tolerance. We observe that the number of realizations decreases slower than linearly from levels 1 to 12.

In Figure 5.11, we see that our dual-weighted estimator of the strong error, $\mathcal{V}_\ell$, gives essentially the same results as the standard Monte Carlo estimator, but with much less computational work. In this case, an accurately empirical estimate of $\mathcal{V}_7$ took almost 48 hours, but the dual-based computation of $\hat{\mathcal{V}}_7$ just took few minutes.

In the simulations, we observed that, as we refine TOL, the optimal number of levels approximately increases logarithmically, which is a desirable feature. We fit the model $L^* = a + b \log(TOL^{-1})$, obtaining b=−2.48 and a=−2.85.

The QQ-plot in the Figure 5.12, computed in the same way as in the previous

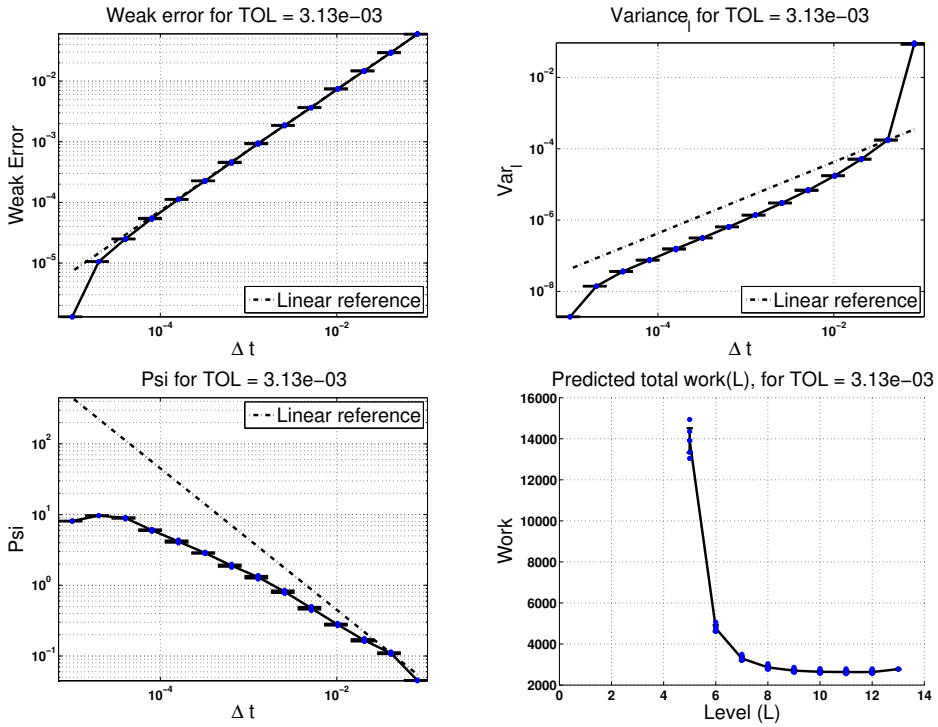

FIG. 5.8. *Upper left: estimated weak error, $\hat{\mathcal{E}}_{I,\ell}$, as a function of the time mesh size, Δt, for the gene transcription and translation model (5.2). Upper right: estimated variance of the difference between two consecutive levels, $\hat{\mathcal{V}}_\ell$, as a function of Δt. Lower left: estimated path work, $\hat{\psi}_\ell$, as a function of Δt. Lower right: estimated total computational work, $\sum_{l=0}^{L} \hat{\psi}_l M_l$, as a function of the level, L.*

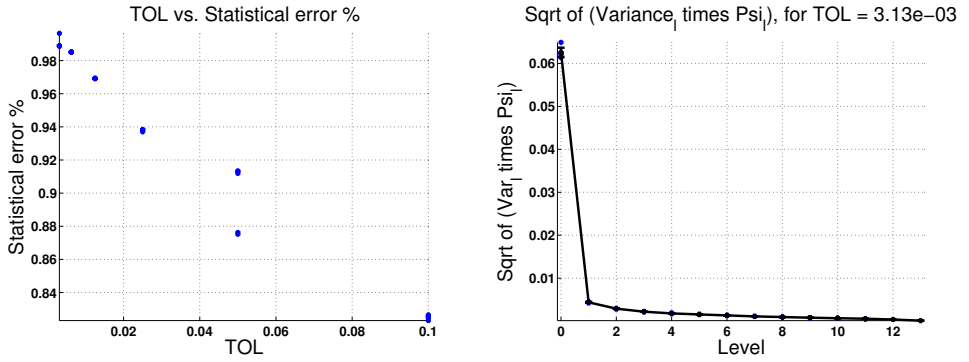
TOL vs. Statistical error %
Statistical error %
TOL
Sqrt of (Variance_l times Psi_l), for TOL = 3.13e−03
Sqrt of (Var_l times Psi_l)
Level

FIG. 5.9. *Left: Percentage of the statistical error over the computational global error, for the gene transcription and translation model (5.2). As mentioned in Section 4, it is well above 0.5 for all the tolerances. Right: $\sqrt{\hat{\mathcal{V}}_\ell \hat{\psi}_\ell}$ as a function of ℓ, for the smallest tolerance, which decreases as the level increases. Observe that the contribution of level 0 is almost equal to the sum of the other levels.*

example, together with a Shapiro-Wilk normality test, shows the validity of the Gaussian assumption for the statistical errors. In the same figure, we also show TOL versus the actual global computational error. It can be seen that the prescribed tolerance

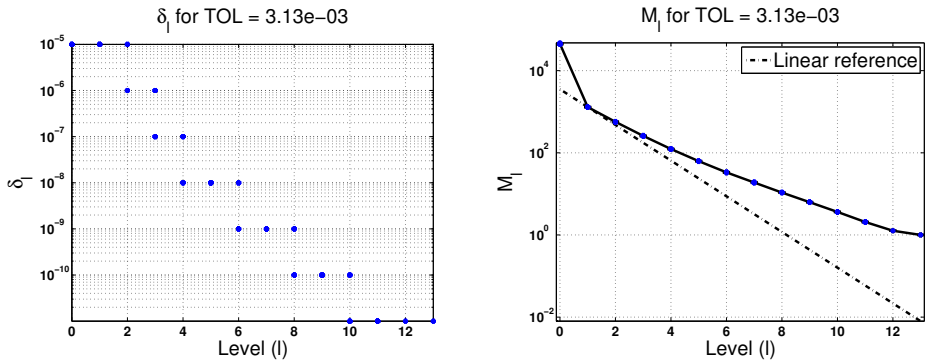

FIG. 5.10. *The one-step exit probability bound, δ_ℓ, and M_ℓ for ℓ=0, 1, ..., L^*, for the smallest tolerance in the gene transcription and translation model (5.2).*

TOL	L^*	Min	Max	$\frac{\hat{W}_{ML}}{\hat{W}_{\rm SSA}}$	Min	Max	$\frac{W_{ML}}{W_{\rm SSA}}$	Min	Max
1.00e-01	3	3	3	0.04	0.04	0.04	0.06	0.05	0.07
5.00e-02	4.6	4	5	0.04	0.03	0.04	0.05	0.05	0.05
2.50e-02	6	6	6	0.03	0.03	0.04	0.05	0.04	0.05
1.25e-02	8	8	8	0.03	0.03	0.03	0.05	0.05	0.06
6.25e-03	10	10	10	0.03	0.03	0.03	0.05	0.04	0.05
3.13e-03	11.4	11	13	0.03	0.03	0.03	0.05	0.04	0.05

TABLE 5.2

Details for the ensemble run of Algorithm 7 for the gene transcription and translation model (5.2).

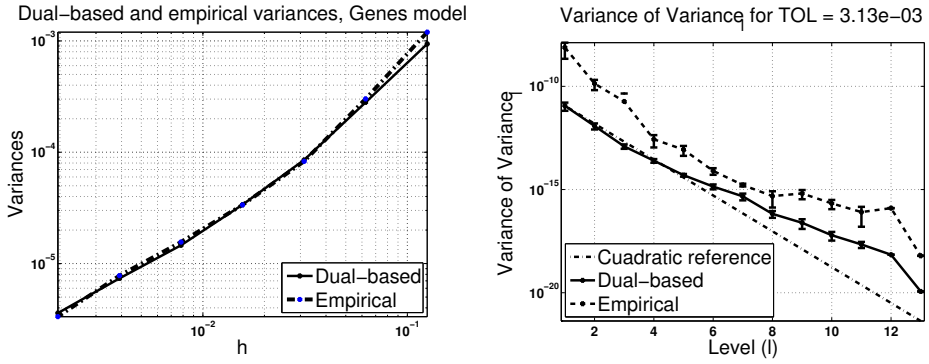

FIG. 5.11. *Left: performance of formula (3.4) as a strong error estimate for the gene transcription and translation model (5.2). Here, h=Δt. Right: estimated variance of $\mathcal{V}_\ell$ with 95% confidence intervals.*

is achieved, except for the second smallest tolerance, with the required confidence of 95%, since C_A=1.96.

MLMC Hybrid-Path Analysis. We now analyze an ensemble of 10^3 independent runs of the multilevel estimator, $\mathcal{M}_L$, for $TOL = 1.25e{-}2$. In this case, $L^* = 8$. In Figures 5.13, 5.14 and 5.15, we show boxplots corresponding to that ensemble. In

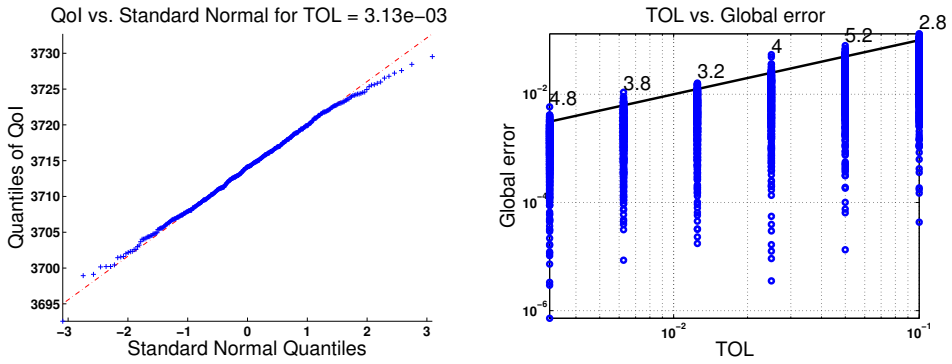

FIG. 5.12. *Left: QQ-plot based on $\mathcal{M}_L$ estimates for the gene transcription and translation model (5.2). Also, we performed a Shapiro-Wilk normality test and we obtained a p-value of* 0.6. *Right: TOL versus the actual global computational error. The numbers above the straight line show the percentage of runs that had errors larger than the required tolerance. We observe that in all cases (except the second for a very small margin) the computational error follows the imposed tolerance with the expected confidence of 95%.*

each one, we indicate the coupling pair (on the x-axis) and the value of δ_ℓ (below the title of the plot). In each figure, the first boxplot starting from the left, corresponds to single-level hybrid simulations at the coarsest level, ℓ=0, with a time mesh of size Δt_0, and with an exit bound for the one-step exit probability, δ_0=1e−5. Next, we show the boxplots corresponding to coupled hybrid paths, at levels ℓ=0 and ℓ=1, generated using time meshes of size Δt_0 and Δt_1, respectively, and exit probability bounds, δ_0 and δ_1, respectively. This is indicated under the boxplots with the symbols 1C and 1F, which stand for 'Coarse' and 'Fine' levels in the first coupling, respectively. We proceed in the same fashion until the final level, L^*. At this point, it is crucial to observe that the probability law for the samples in the boxplots indicated by kF and $(k+1)C$ should be the same for any $k \in \{0, 1, \ldots, L^*-1\}$ (in the single-level case, we interpret the symbol 0 as 0F). This is because both samples are generated using the same time mesh of size Δt_k and the same one-step exit probability bound, δ_k, see Remark 2.1.

Figure 5.13 shows the total proportion of Chernoff tau-leap steps over the total number of tau-leap steps. Here, we understand that a Chernoff tau-leap step is taken when the size of τ_{Ch} (see Section 1.5) is strictly smaller than the distance from the current time to the next mesh point and, therefore, the Chernoff bound is acting as an actual constraint for the size of the tau-leap step. We can see how the Chernoff steps are present in the first levels but not in the final ones, where exact steps are preferred according to our computational work criterion. We observe a small increase of the proportion of the number of Chernoff steps from levels 1F/2C to levels 3F/4C (strictly speaking, a shift in the median and the third quartile). This is due to consecutive refinements in the values of δ, from 1e−5 to 1e−7, producing smaller and smaller values of τ_{Ch}. This is also because the Chernoff step size is, for some time points, still smaller than the grid size and also because the cost of reaching the time horizon using Chernoff steps is still preferred over the cost of using exact steps. The abundance of outliers at all levels up to ℓ=3 indicates that the Chernoff bound is actively controlling the size of the tau-leap steps.

Figures 5.14 and 5.15 show the total count of tau-leap and exact steps, respectively. These plots are intended as diagnostic plots with two main objectives: i)

checking the telescoping sum property as stated in Remark 2.1, and ii) understanding the 'blending' phenomena in our simulated hybrid paths, that is, the presence of both methods, tau-leap and exact. It could be useful to think in terms of the domain of each method: given a time mesh, Δt, and a value of the one-step exit probability bound, δ, we could decompose the interval $[0,T]$ into two domains, I_{TL} and I_{MNRM}, for the tau-leap and exact methods, respectively. The domain, I_{MNRM}, should be monotonically increasing with refinements of the time mesh and δ, since when the size of the time mesh, Δt_ℓ or δ_ℓ, goes to zero, the expected number of tau-leap steps also goes to zero, see [19], Appendix A. As a consequence, we expect the total count of exact paths to be a monotonically increasing function of the level, ℓ. On the other hand, the domain I_{TL} decreases, but, since the size of the time mesh halves form by passing from one level to the next one, we expect to see also an increasing number of tau-leap steps, at least for no very deep levels. The blending effect of the hybrid decision rules in Algorithm 3 are depicted in Figure 5.16, where the proportion of the tau-leap steps over the total number of steps is shown for levels $\ell \in \{0, 5, 8\}$. In the left panel, we can see that the number of tau-leap steps dominates except close to the origin, where the coarse time-mesh is finer. Remember that in our methodology, our initial mesh can be nonuniform. We then see how the domain, I_{MNRM}, increases until it occupies almost 80% of the time interval $[0,T]$.

REMARK 5.1. *The savings in computational work when generating Poisson random variables heavily depend on MATLAB's performance capabilities. For example, we do not generate the random variates in batches, as in [1], and that could have an impact on the results. In fact, we should expect better results from our method if we implement our algorithms in more performance-oriented languages or if we sample Poisson random variables in batches.*

REMARK 5.2. *(Level 0 time mesh) In this example, we use an adaptive mesh at level 0. This is because this example is mildly stiff. Using a uniform time mesh at level 0 imposes a small time step size requirement for all time which is not needed. Moreover, this issue is propagated to the finer levels. In all our numerical examples, at level 0, we use the coarsest possible time mesh such that the Forward Euler method is numerically stable.*

6. Conclusions. In this work, we developed a multilevel Monte Carlo version for the single-level hybrid Chernoff tau-leap algorithm presented in [19]. We showed that the computational complexity of this method is of order $\mathcal{O}\left(TOL^{-2}\right)$ and, therefore, that it can be seen as a variance reduction of the SSA method, which has the same complexity. This represents an important advantage of the hybrid tau-leap with respect to the pure tau-leap in the multilevel context. In our numerical examples, we obtained substantial gains with respect to both the SSA and the single-level hybrid Chernoff tau-leap. The present approach, like the one in [19], also provides an approximation of $\mathrm{E}\left[g(X(T))\right]$ with prescribed accuracy and confidence level, with nearly optimal computational work. For reaching this optimality, we derived novel formulas based on dual-weighted residual estimations for computing the variance of the difference of the observables between two consecutive levels in coupled hybrid paths and also the bias of the deepest level (see (3.3) and (3.4)). These formulas are particularly relevant in the present context of Stochastic Reaction Networks due to the fact that alternative standard sample estimators become too costly at deep levels because of the presence of large kurtosis.

Future extensions may involve better hybridization techniques as well as implicit and higher-order versions of the hybrid MLMC.

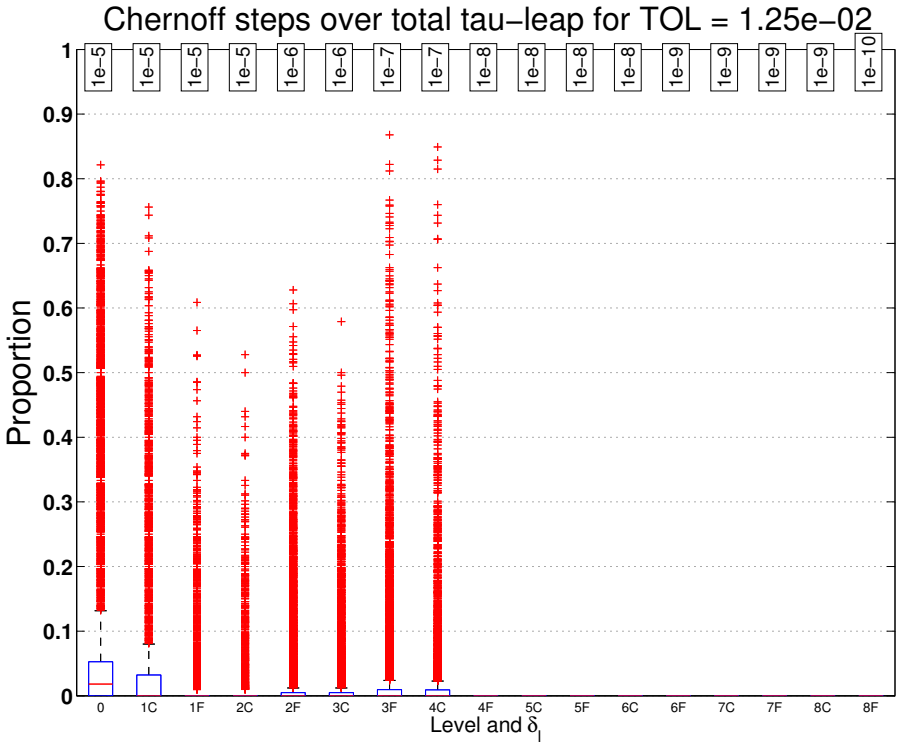

FIG. 5.13. *Proportion of the number of Chernoff tau-leap steps over the total number of tau-leap steps for the gene transcription and translation model (5.2). In the x-axis, we show the corresponding level (starting from level 0) and, subsequently, the coarse (C) and fine (F) level. Below the title, we show the corresponding* δ_ℓ *of each level. We observe a small increase in the proportion of the number of Chernoff steps from levels 1F/2C to levels 3F/4C (strictly speaking, a shift in the median and the third quartile). This is due to consecutive refinements in the values of* δ*, from* $1e{-}5$ *to* $1e{-}7$*, producing smaller and smaller values of* τ_{Ch}*.*

Acknowledgments. The authors would like to thank two anonymous reviewers for their constructive comments that helped us to improve our manuscript. We also would like to thank Prof. Mike Giles for very enlightening discussions. The authors are members of the KAUST SRI Center for Uncertainty Quantification in the Computer, Electrical and Mathematical Sciences and Engineering Division at King Abdullah University of Science and Technology (KAUST). This work was supported by KAUST.

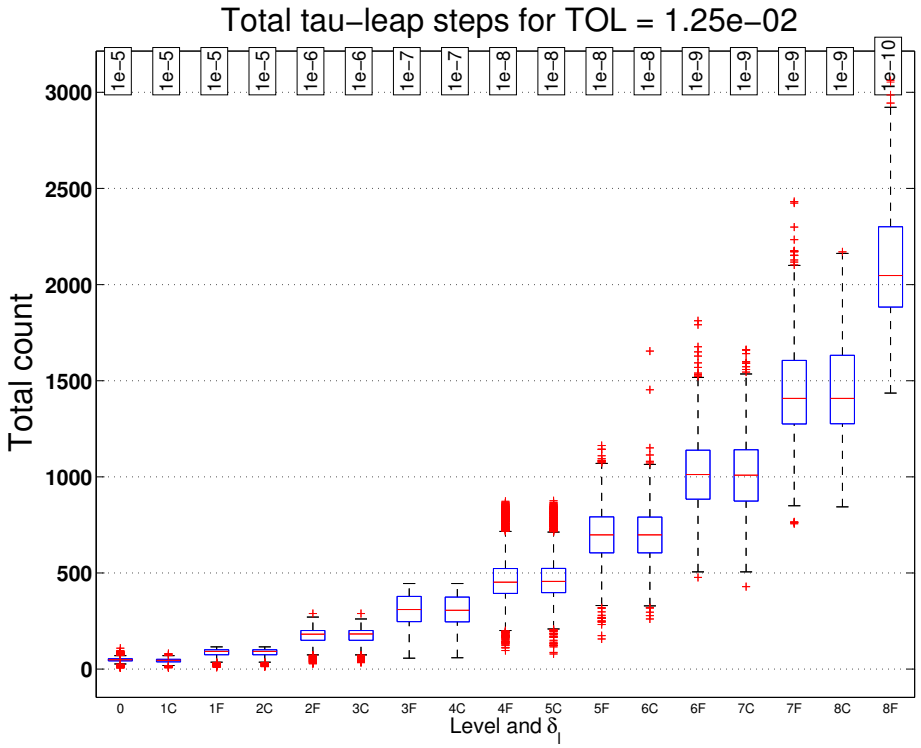

FIG. 5.14. *Total number of tau-leap steps per path for the gene transcription and translation model (5.2). In the x-axis, we show the corresponding pairings of two consecutive levels (starting from level 0) and, subsequently, the coarse (C) and fine (F) meshes for two consecutive levels. Below the title, we show the corresponding δ_ℓ of each level. The domain I_{TL} of the tau-leap method decreases with refinements, but, since the size of time mesh halves form by passing from one level to the next one, we see an increasing number of tau-leap steps until, at a certain level, there are no more tau-leap steps due to the relative computational cost of the tau-leap method.*

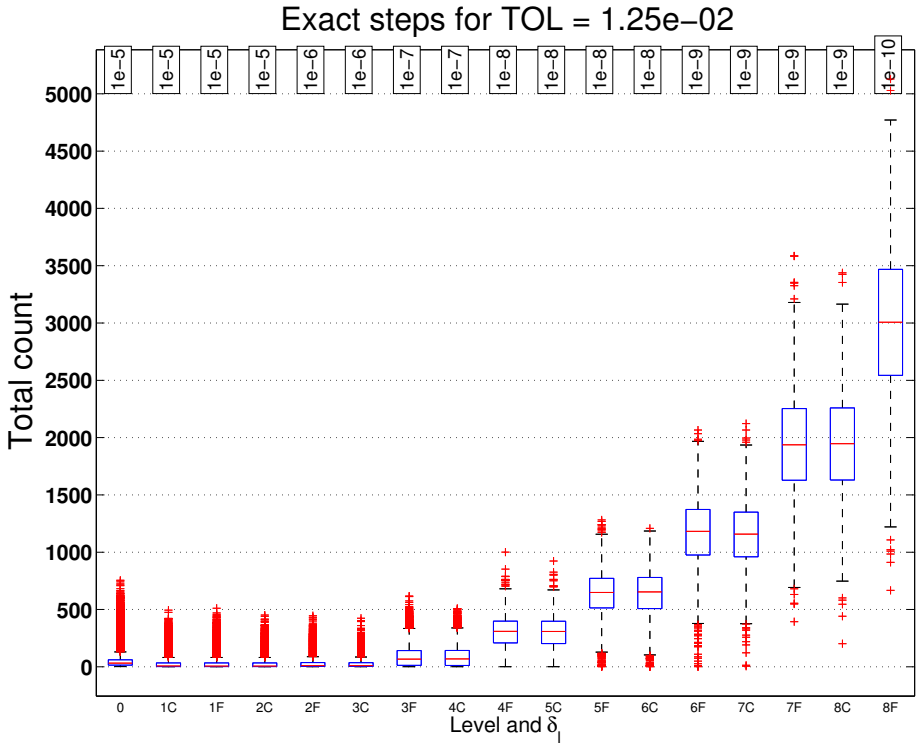

FIG. 5.15. *Total number of exact steps per path for the gene transcription and translation model (5.2). In the x-axis, we show the corresponding pairings of two consecutive levels (starting from level 0) and, subsequently, the coarse (C) and fine (F) meshes for two consecutive levels. Below the title, we show the corresponding δ_ℓ of each level. The domain I_{MNRM} of the exact method is monotonically increasing with refinements of the time mesh and the one-step exit probability bound. As a consequence, we expect the total count of exact paths to be a monotonically increasing function of the level, ℓ.*

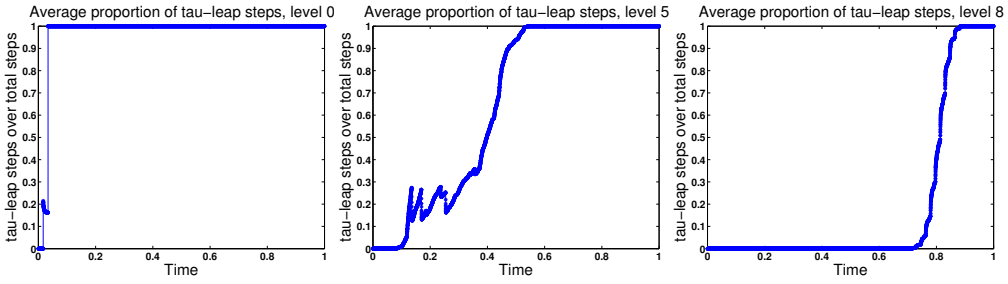
Average proportion of tau−leap steps, level 0
Average proportion of tau−leap steps, level 5
Average proportion of tau−leap steps, level 8
tau−leap steps over total steps
Time

FIG. 5.16. *This figure depicts the 'blending' effect produced by our hybrid path-simulation algorithm. Here, we can see the proportion of the tau-leap steps adaptively taken based on expected work optimization. We see how the presence of the tau-leap decreases when we move to the deepest levels. We observe that, for the chosen tolerance, to couple with an exact path at the last level is not optimal.*

REFERENCES

[1] D. Anderson and D. Higham. Multilevel Monte Carlo for continuous Markov chains, with applications in biochemical kinetics. *SIAM Multiscal Model. Simul.*, 10(1), 2012.

[2] D. Anderson, D. Higham, and Y. Sun. Complexity of multilevel monte carlo tau-leaping. *arXiv:1310.2676v1*, 2013.

[3] D. F. Anderson. A modified next reaction method for simulating chemical systems with time dependent propensities and delays. *The Journal of Chemical Physics*, 127(21), 2007.

[4] J. P. Aparicio and H. Solari. Population dynamics: Poisson approximation and its relation to the langevin processs. *Physical Reveiw Letters*, 86(18):4183–4186, Apr. 2001.

[5] C. Bierig and A. Chernov. Convergence analysis of multilevel variance estimators in multilevel monte carlo methods and application for random obstacle problems. Preprint 1309, Institute for Numerical Simulation, University of Bonn, 2013. Submitted.

[6] N. Collier, A.-L. Haji-Ali, F. Nobile, E. von Schwerin, and R. Tempone. A continuation multilevel monte carlo algorithm. *Mathematics Institute of Computational Science and Engineering, Technical report Nr. 10.2014, EPFL*, 2014.

[7] S. N. Ethier and T. G. Kurtz. *Markov Processes: Characterization and Convergence (Wiley Series in Probability and Statistics)*. Wiley-Interscience, 2nd edition, 9 2005.

[8] M. A. Gibson and J. Bruck. Efficient exact stochastic simulation of chemical systems with many species and many channels. *The Journal of Physical Chemistry A*, 104(9):1876–1889, 2000.

[9] M. Giles. Multi-level Monte Carlo path simulation. *Operations Research*, 53(3):607–617, 2008.

[10] D. T. Gillespie. A general method for numerically simulating the stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22:403–434, 1976.

[11] D. T. Gillespie. Approximate accelerated stochastic simulation of chemically reacting systems. *Journal of Chemical Physics*, 115:1716–1733, July 2001.

[12] S. Heinrich. Monte carlo complexity of global solution of integral equations. *Journal of Complexity*, 14(2):151–175, 1998.

[13] S. Heinrich. Multilevel monte carlo methods. In *Large-Scale Scientific Computing*, volume 2179 of *Lecture Notes in Computer Science*, pages 58–67. Springer Berlin Heidelberg, 2001.

[14] J. Karlsson, M. Katsoulakis, A. Szepessy, and R. Tempone. Automatic Weak Global Error Control for the Tau-Leap Method. *preprint arXiv:1004.2948v3*, pages 1–22, 2010.

[15] J. Karlsson and R. Tempone. Towards automatic global error control: Computable weak error expansion for the tau-leap method. *Monte Carlo Methods and Applications*, 17(3):233–278, March 2011.

[16] T. G. Kurtz. Representation and approximation of counting processes. *Adv. Filtering OptimalStochastic Control 42*, 1982.

[17] T. Li. Analysis of explicit tau-leaping schemes for simulating chemically reacting systems. *Multiscale Modeling and Simulation*, 6(2):417–436, 2007.

[18] D. G. Luenberger and Y. Ye. *Linear and Nonlinear Programming (International Series in Operations Research and Management Science)*. Springer, 2010.

[19] A. Moraes, R. Tempone, and P. Vilanova. Hybrid chernoff tau-leap. *Multiscale Modeling & Simulation*, 12(2):581–615, 2014.

[20] S. S. Shapiro and M. B. Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611, Dec. 1965.

[21] A. Speight. A multilevel approach to control variates. *Journal of Computational Finance*, 12:1–25, 2009.

Algorithm 3 Coupled hybrid path. Inputs: the initial state, $X(0)$, the final time, T, the propensity functions, $a{=}(a_j)_{j=1}^J$, the stoichiometric vectors, $\nu{=}(\nu_j)_{j=1}^J$, two one-step exit probability bounds; one for the coarse level, $\bar{\delta}$, and another for the fine level, $\bar{\bar{\delta}}$, and two time meshes, one coarse $(t_k)_{k=0}^K$, such that $t_K{=}T$ and a finer one, $(s_l)_{l=0}^{K'}$, such that $s_0{=}t_0$, $s_{K'}{=}t_K$, and $(t_k)_{k=0}^K \subset (s_l)_{l=0}^{K'}$. Outputs: a sequence of states evaluated at the coarse grid, $(\bar{X}(t_k))_{k=0}^K \subset \mathbb{Z}_+^d$, such that $t_K \leq T$, a sequence of states evaluated at the fine grid $(\bar{\bar{X}}(s_l))_{l=0}^{K'} \subset \mathbb{Z}_+^d$, such that $\bar{X}(t_K) \in \mathbb{Z}_+^d$ or $\bar{\bar{X}}(s_{K'}) \in \mathbb{Z}_+^d$. If $t_K < T$, both paths exited the $\mathbb{Z}_+^d$ lattice before the final time, T. It also returns the number of times the tau-leap method was successfully applied at the fine level and at the coarse level, and the number of exact steps at the fine level and at the coarse level. For the sake of simplicity, we omit sentences involving the recording of $\bar{X}(t_k)$ and $\bar{\bar{X}}(s_l)$ from the current state variables $\bar{X}$ and $\bar{\bar{X}}$, respectively, the counting of the number of steps, and the return sentence.

$t \leftarrow 0$, $\bar{X} \leftarrow X(0)$, $\bar{\bar{X}} \leftarrow X(0)$
$\bar{t} \leftarrow$ next grid point in the coarse grid larger than t
$(\bar{H}, \bar{m}, \bar{a}) \leftarrow$ Algorithm 4 with $(\bar{X}, t, \bar{t}, T, \bar{\delta}, a)$
$\bar{\bar{t}} \leftarrow$ next grid point in the fine grid larger than t
$(\bar{\bar{H}}, \bar{\bar{m}}, \bar{\bar{a}}) \leftarrow$ Algorithm 4 with $(\bar{\bar{X}}, t, \bar{\bar{t}}, T, \bar{\bar{\delta}}, a)$
while $t < T$ **do**
  $H \leftarrow \min\{\bar{H}, \bar{\bar{H}}\}$
  if $\bar{m} = \text{TL}$ **and** $\bar{\bar{m}} = \text{TL}$ **then**
    $S \leftarrow$ Algorithm 6 with $(\bar{a}, \bar{\bar{a}})$
    $\Lambda \leftarrow \mathcal{P}(S{\cdot}(H{-}t))$ (generate Poisson random variates)
    $\bar{\bar{X}} \leftarrow \bar{\bar{X}} + (\Lambda_1{+}\Lambda_2)\nu$
    $\bar{X} \leftarrow \bar{X} + (\Lambda_1{+}\Lambda_3)\nu$
    $t \leftarrow H$
  else
    Initialize internal clocks R, P if needed (see Algorithm 1)
    while $t < H$ **do**
      if $\bar{\bar{m}} = \text{MNRM}$ **then**
        $\bar{\bar{a}} \leftarrow a(\bar{\bar{X}})$
      end if
      if $\bar{m} = \text{MNRM}$ **then**
        $\bar{a} \leftarrow a(\bar{X})$
      end if
      $S \leftarrow$ Algorithm 6 with $(\bar{a}, \bar{\bar{a}})$
      $(t, \bar{X}, \bar{\bar{X}}, R, P) \leftarrow$ Algorithm 5 with $(t, H, \bar{X}, \bar{\bar{X}}, R, P, S)$
    end while
  end if
  if $t < T$ **then**
    if $H = \bar{H}$ **then**
      $\bar{t} \leftarrow$ next grid point in the coarse grid larger than t
      $(\bar{H}, \bar{m}, \bar{a}) \leftarrow$ Algorithm 4 with $(\bar{X}, t, \bar{t}, T, \bar{\delta}, a)$
    end if
    if $H = \bar{\bar{H}}$ **then**
      $\bar{\bar{t}} \leftarrow$ next grid point in the fine grid larger than t
      $(\bar{\bar{H}}, \bar{\bar{m}}, \bar{\bar{a}}) \leftarrow$ Algorithm 4 with $(\bar{\bar{X}}, t, \bar{\bar{t}}, T, \bar{\bar{\delta}}, a)$
    end if
  end if
end while

Algorithm 4 Compute next time horizon. Inputs: the current state, $\tilde{X}$, the current time, t, the next grid point, $\tilde{t}$, the final time, T, the one step exit probability bound, $\tilde{\delta}$, and the propensity functions, $a{=}(a_j)_{j=1}^J$. Outputs: the next horizon H, the selected method m, current propensity values $\tilde{a}$.

$\tilde{a} \leftarrow a(\tilde{X})$
$(m, \tilde{\tau}) \leftarrow$ Algorithm 2 with $(\tilde{X}, t, \tilde{a}, \tilde{\delta}, \tilde{t})$
if $m = \text{TL}$ **then**
  $H \leftarrow \min\{\tilde{t}, t{+}\tilde{\tau}, T\}$
else
  $H \leftarrow \min\{t{+}\tilde{\tau}, T\}$
end if
return $(H, m, \tilde{a})$

Algorithm 5 Auxiliary function used in Algorithm 3. Inputs: the current time, t, the current time horizon, H, the current system state at coarser level, $\bar{X}$, and finer level, $\bar{\bar{X}}$, the internal clocks R_i, P_i, $i{=}1,2,3$, and the values, S_i, $i{=}1,2,3$ (see Section 1.2 for more information on these values). Outputs: updated time, t, updated system states, $\bar{X}$, $\bar{\bar{X}}$, and updated internal clocks R_i, P_i, $i{=}1,2,3$.

$\Delta t_i \leftarrow (P_i - R_i)/S_i$, for $i{=}1,2,3$
$\Delta \leftarrow \min_i\{\Delta t_i\}$
$\mu \leftarrow \text{argmin}_i\{\Delta t_i\}$
if $t + \Delta > H$ **then**
  $R \leftarrow R + S{\cdot}(H{-}t)$
  $t \leftarrow H$
else
  update $\bar{X}$ and $\bar{\bar{X}}$
  $R \leftarrow R + S\Delta$
  $r \leftarrow \text{uniform}(0,1)$
  $P_\mu \leftarrow P_\mu + \log(1/r)$
  $t \leftarrow t + \Delta$
end if
return $(t, \bar{X}, \bar{\bar{X}}, R, P)$

Algorithm 6 Auxiliary function used in Algorithm 3. Inputs: the propensity values at the coarse and fine grid, $\bar{a}, \bar{\bar{a}}$. Output: S_i, $i{=}1,2,3$.

$S_1 \leftarrow \min(\bar{a}, \bar{\bar{a}})$
$S_2 \leftarrow \bar{a} - S_1$
$S_3 \leftarrow \bar{\bar{a}} - S_1$
return S

Algorithm 7 Multilevel calibration and error estimation. Inputs: same as Algorithm 3 plus the observable, g, and the prescribed tolerance, $TOL>0$. Outputs: $(M_\ell)_{\ell=0}^{L}$, $(\delta_\ell)_{\ell=0}^{L}$, $((t_{n,\ell})_{n=0}^{N_\ell})_{\ell=0}^{L}$, the estimated computational work of the multilevel estimator, $\hat{W}_{\mathrm{ML}}$, and the estimated computational work of the SSA method, $\hat{W}_{\mathrm{SSA}}$. We denote by $g_l \equiv g(\bar{X}_l(T;\bar{\omega}))$, and $g_{l+1}-g_l \equiv g(\bar{X}_{l+1}(T;\bar{\omega}))-g(\bar{X}_l(T;\bar{\omega}))$. Here, C^* is the unitary cost of a pure SSA step, and c is the factor of refinement of δ (in our experiments c=10). See also Remark 4.1 regarding the estimators of $\mathrm{Var}\left[g(X(T))\right]$ and $\mathrm{E}\left[g(X(T))\right]$, and Remark 4.3.

$l \leftarrow 0$, $\delta_l \leftarrow 0.01$, $\hat{W}_{ML}^{(a)} \leftarrow \infty$
Set initial meshes $(t_k)_{k=0}^{K}$ and $(s_l)_{l=0}^{K'}$
fin-delta $\leftarrow$ **false**
while not fin-delta **do**
  $(\hat{\psi}_0, \mathcal{S}^2\left(g_l;\cdot\right), \mathcal{A}\left(\{g_l, \mathcal{E}_I, N_{\mathrm{SSA}^*}, N_{\mathrm{TL}}\};\cdot\right)) \leftarrow$ Algorithm 12
  if $\hat{\mathcal{V}}_l(1-\delta_l\mathcal{A}\left(N_{\mathrm{TL}};\cdot\right))^2 \geq 2\mathcal{S}^2\left(g_l;\cdot\right)\delta_l\mathcal{A}\left(N_{\mathrm{TL}};\cdot\right)$ **and** $\delta_l\mathcal{A}\left(N_{\mathrm{TL}};\cdot\right) < 0.1$ **then**
    fin-delta $\leftarrow$ **true**
    Refine δ_l by a factor of c
  end if
end while
$\delta_{l+1} \leftarrow \delta_l$
fin $\leftarrow$ **false**
while not fin **do**
  fin-delta $\leftarrow$ **false**
  while not fin-delta **do**
    $(\hat{\psi}_{l+1}, \hat{\mathcal{V}}_{l+1}, \mathcal{A}\left(\{g_{l+1}, N_{\mathrm{SSA}^*}, \mathcal{E}_I, N_{TL,l+1}\};\cdot\right), \mathcal{S}^2\left(g_{l+1};\cdot\right)) \leftarrow$ Algorithm 8
    if $\hat{\mathcal{V}}_{l+1}(1-\delta_{l+1}\mathcal{A}\left(N_{TL,l+1};\cdot\right))^2 \geq 2\mathcal{S}^2\left(g_{l+1};\cdot\right)\delta_{l+1}\mathcal{A}\left(N_{TL,l+1};\cdot\right)$
    and $\delta_{l+1}\mathcal{A}\left(N_{TL,l+1};\cdot\right) < 0.1$ **then**
      fin-delta $\leftarrow$ **true**
      $\delta_l \leftarrow \delta_{l+1}$
    else
      Refine δ_{l+1} by a factor of c
    end if
  end while
  $M_{\mathrm{SSA}} \leftarrow C_A^2\mathcal{S}^2\left(g_{l+1};\cdot\right)/TOL^2$
  $\hat{W}_{\mathrm{SSA}} \leftarrow C^* M_{\mathrm{SSA}}\mathcal{A}\left(N_{\mathrm{SSA}^*};\cdot\right)$
  if $\hat{\mathcal{E}}_I < TOL-TOL^2$ **then**
    $(M_\ell)_{\ell=0}^{l+1} \leftarrow$ Algorithm 11 with $((\hat{\psi}_\ell)_{\ell=0}^{l+1}, (\hat{\mathcal{V}}_\ell)_{\ell=0}^{l+1}, TOL, \hat{\mathcal{E}}_I)$
    $\hat{W}_{\mathrm{ML}} \leftarrow \sum_{\ell=0}^{l+1}\hat{\psi}_\ell M_\ell$
  else
    $\hat{W}_{\mathrm{ML}} \leftarrow \infty$
  end if
  if $(\hat{W}_{\mathrm{ML}}^{(a)} > \hat{W}_{\mathrm{ML}}$ **or** $\hat{\mathcal{E}}_I > TOL-TOL^2)$ **and** $\mathcal{A}\left(N_{TL,l+1};\cdot\right) > 0$ **then**
    $l \leftarrow l+1$
    $\hat{W}_{\mathrm{ML}}^{(a)} \leftarrow \hat{W}_{\mathrm{ML}}$
    Refine meshes $(t_k)_{k=0}^{K}$ and $(s_l)_{l=0}^{K'}$
  else
    fin $\leftarrow \delta_{l+1}\mathcal{A}\left(N_{\mathrm{TL},l+1};\cdot\right)\mathcal{A}\left(g_{l+1};\cdot\right) \leq TOL^2$
    if not fin **then**
      $\delta_{l+1} \leftarrow c^{\lfloor(log_c(TOL^2/(\mathcal{A}\left(g_{l+1};\cdot\right)\cdot\mathcal{A}\left(N_{\mathrm{TL},l+1};\cdot\right))\rfloor}$
      while not fin **do**
        $(\mathcal{A}\left(\{g_{l+1}, N_{TL,l+1}\};\cdot\right)) \leftarrow$ Algorithm 8
        fin $\leftarrow \delta_{l+1}\mathcal{A}\left(N_{\mathrm{TL},l+1};\cdot\right)\mathcal{A}\left(g_{l+1};\cdot\right) \leq TOL^2$
        if not fin **then**
          Refine δ_{l+1} by a factor of c
        end if
      end while
    end if
    if $\mathcal{A}\left(N_{\mathrm{TL},l+1};\cdot\right) = 0$ **then**
      $l \leftarrow l+1$
      $(M_\ell)_{\ell=0}^{l} \leftarrow$ Algorithm 11 with $((\hat{\psi}_\ell)_{\ell=0}^{l}, (\hat{\mathcal{V}}_\ell)_{\ell=0}^{l}, TOL, 0)$
      $\hat{W}_{\mathrm{ML}} \leftarrow \sum_{\ell=0}^{l}\hat{\psi}_\ell M_\ell$
    end if
  end if
end while

Algorithm 8 Auxiliary function for Algorithm 7. Inputs: same as Algorithm 3. Outputs: the estimated runtime of the coupled path, $\hat{\psi}$, an estimate of $\mathrm{Var}\left[g(\bar{X}(T))-g(\bar{\bar{X}}(T))\right]$, $\hat{\mathcal{V}}$, an estimate of $\mathrm{E}\left[g(X(T))\right]$, $\mathcal{A}\left(g(\bar{\bar{X}}(T);\cdot\right)$, an estimate of the expected number of steps needed by the SSA method, $\mathcal{A}\left(N_{\mathrm{SSA}^*};\cdot\right)$, an estimate of $\mathrm{E}\left[\mathcal{E}_I\right]$, $\mathcal{A}\left(\mathcal{E}_I;\cdot\right)$, an estimate of the expected number of tau-leap steps taken at the fine level, $\mathcal{A}\left(N_{\mathrm{TL}};\cdot\right)$, and an estimate of $\mathrm{Var}\left[g(X(T))\right]$, $\mathcal{S}^2\left(g(\bar{\bar{X}}(T);\cdot\right)$. Here, $\bar{\bar{X}}(t)$ refers to the approximated process using a finer grid than the approximated process, $\bar{X}(t)$. Moreover, $(\bar{X}(t)$,$\bar{\bar{X}}(t))$ are two coupled paths. Here, $\mathbf{1}_{TL}(k)=1$ if and only if the decision at time t_k was tau-leap. Set appropriate values for M_0 and CV_0. For the sake of simplicity, we omit the arguments of the algorithms when there is no risk of confusion. See also Remark 4.1 regarding the estimators of $\mathrm{Var}\left[g(X(T))\right]$ and $\mathrm{E}\left[g(X(T))\right]$.

$M \leftarrow M_0$, $cv \leftarrow \infty$, $M_f \leftarrow 0$
while $cv > CV_0$ **do**
  for $m \leftarrow 1$ to M **do**
    Generate two coupled paths: $(\bar{X}(s_l;\bar{\bar{\omega}}_m))_{l=0}^{K'}, (\bar{\bar{X}}(s_l;\bar{\bar{\omega}}_m))_{l=0}^{K'}, \leftarrow$ Algorithm 3
    if the path does not exit $\mathbb{Z}_+^d$ **then**
      $M_f \leftarrow M_f + 1$
      $(S_e(\bar{\bar{\omega}}_m), S_v(\bar{\bar{\omega}}_m)) \leftarrow$ Algorithm 10 with $(\bar{X}(t_k;\bar{\bar{\omega}}_m))_{k=0}^{K}$
      $\mathcal{E}_I(\bar{\bar{\omega}}_m) \leftarrow$ Algorithm 9 with $(\bar{\bar{X}}(s_l;\bar{\bar{\omega}}_m))_{l=0}^{K'}$
      Estimate $N_{\mathrm{SSA}^*}(\bar{\bar{\omega}}_m)$, using $\int_0^T a_0(\bar{\bar{X}}(s))ds$ (see [19])
      $C_{Poi}(\bar{\bar{\omega}}_m) \leftarrow \sum_{j=1}^{J}\sum_{l=0}^{K'} C_P(a_j(\bar{X}(s_l))(s_{l+1}-s_l)\mathbf{1}_{TL}(l)$
        $+ \sum_{j=1}^{J}\sum_{l=0}^{K'} C_P(a_j(\bar{\bar{X}}(s_l))(s_{l+1}-s_l)\mathbf{1}_{TL}(l)$
      Compute $N_{K1}^{(c)}$, $N_{K2}^{(c)}$, $N_{\mathrm{TL}}^{(c)}$, and N_{TL}
    end if
  end for
  $\hat{\mathcal{V}} \leftarrow \mathcal{S}^2\left(S_e;M_f\right) + \mathcal{A}\left(S_v;M_f\right)$
  Compute the coefficient of variation $cv_{\mathcal{V}}$ and $cv_{\mathcal{E}_I}$ of $\hat{\mathcal{V}}$ and $\mathcal{A}\left(\mathcal{E}_I;\cdot\right)$, respectively.
  $cv \leftarrow \max\{cv_{\mathcal{V}}, cv_{\mathcal{E}_I}\}$
  $\hat{\psi} \leftarrow C_1\mathcal{A}\left(N_{K1}^{(c)};M_f\right) + C_2\mathcal{A}\left(N_{K2}^{(c)};M_f\right) + C_3\mathcal{A}\left(N_{\mathrm{TL}}^{(c)};M_f\right) + \mathcal{A}\left(C_{Poi};M_f\right)$
  $M \leftarrow 2M$
end while
return $\left(\hat{\psi}, \hat{\mathcal{V}}, \mathcal{A}\left(\{g(\bar{\bar{X}}(T)), N_{\mathrm{SSA}^*}, \mathcal{E}_I, N_{\mathrm{TL}}\};M_f\right), \mathcal{S}^2\left(g(\bar{\bar{X}}(T));M_f\right)\right)$

Algorithm 9 Compute the discretization error of a given approximated path. Inputs: $(\bar{X}(t_k))_{k=0}^K$. Here, $\mathbf{1}_{TL}(k) = 1$ if and only if the decision at time t_k was tau-leap, and Id is the $d \times d$ identity matrix Output: $\mathcal{E}_I$. Notes: $x_k \equiv \bar{X}(t_k)$.

$\mathcal{E}_I \leftarrow 0$
Compute $\varphi_K \leftarrow \nabla g(x_k)$
for $k \leftarrow K-1$ to 1 **do**
  $\Delta t_k \leftarrow t_{k+1} - t_k$
  Compute $\mathbb{J}_a = [\partial_i a_j(x_k)]_{j,i}$
  $\varphi_k \leftarrow \left(Id + \Delta t_k \, \mathbb{J}_a^T \, \nu^T\right) \varphi_{k+1}$
  $\Delta a_k \leftarrow a(x_{k+1}) - a(x_k)$
  $\mathcal{E}_I \leftarrow \mathcal{E}_I + \frac{\Delta t_k}{2} (\Delta a_k \, \mathbf{1}_{TL}(k) \, \nu^T) \, \varphi_k$
end for
return $\mathcal{E}_I$

Algorithm 10 Compute $S_e \equiv S_e(\bar{\omega})$ and $S_v \equiv S_v(\bar{\omega})$ defined in (3.12). Inputs: $(\bar{X}(t_k))_{k=0}^K$ and a positive constant c. Outputs: S_e and S_v. Notes: if a is a vector, then, $\mathrm{diag}(a)$ is a diagonal matrix with main diagonal a. Here, $\mathbf{1}_{TL}(k) = 1$ if and only if the decision at time t_k was tau-leap, Id is the $d \times d$ identity matrix, $x_k \equiv \bar{X}(t_k)$, and $\Phi(x)$ is the cumulative distribution function of a Gaussian random variable.

$S_e \leftarrow 0$
$S_v \leftarrow 0$
Compute $\varphi_K \leftarrow \nabla g(x_k)$
for $k \leftarrow K-1$ to 1 **do**
  $\Delta t_k \leftarrow t_{k+1} - t_k$
  Compute $\mathbb{J}_a = [\partial_i a_j(x_k)]_{j,i}$
  $\varphi_k \leftarrow \left(Id + \Delta t_k \, \mathbb{J}_a^T \, \nu^T\right) \varphi_{k+1}$
  $\nu_\varphi \leftarrow \nu^T \varphi_k$
  $\nu_a \leftarrow (\mathbb{J}_a \, \nu)^T$
  $\mu_j \leftarrow \frac{\Delta t_k}{2} \sum_i (\nabla a_j(x_k) \cdot \nu_i) \, a_i(x_k)$
  $\bar{\mu}_j \leftarrow \frac{\Delta t_k}{2} \sum_i |(\nabla a_j(x_k) \cdot \nu_i)| \, a_i(x_k)$
  $\sigma_j^2 \leftarrow \frac{\Delta t_k}{2} \sum_i (\nabla a_j(x_k) \cdot \nu_i)^2 \, a_i(x_k)$
  $S_e \leftarrow S_e + \mathbf{1}_{\mathrm{TL}}(k) \frac{\Delta t_k}{2} \mu \, \nu_\varphi$
  $\mathrm{aux}_1 \leftarrow \frac{(\Delta t_k)^3}{8} (\nu_a \nu_\varphi)^T \mathrm{diag}(a) (\nu_a \nu_\varphi)$
  $\mathrm{aux}_2 \leftarrow \frac{\Delta t_k}{2} \sum_j (\varphi_k \cdot \nu_j)^2 \mathbf{1}_{\left\{\frac{\Delta t_k}{2} a_j(x_k) > c\right\}} \left(\mu_j (1 - 2\Phi(-\frac{\mu_j}{\sigma_j})) + \sqrt{\frac{2}{\pi}} \, \sigma_j \, \exp(-\frac{1}{2}(\frac{\mu_j}{\sigma_j})^2)\right)$
  $\mathrm{aux}_3 \leftarrow \frac{\Delta t_k}{2} \sum_j (\varphi_k \cdot \nu_j)^2 \mathbf{1}_{\left\{\frac{\Delta t_k}{2} a_j(x_k) < c\right\}} \min\left\{\bar{\mu}_j, \sqrt{\mu_j^2 + \sigma_j^2}\right\}$
  $S_v \leftarrow S_v + \mathbf{1}_{\mathrm{TL}}(k) (\mathrm{aux}_1 + \mathrm{aux}_2 + \mathrm{aux}_3)$
end for
return (S_e, S_v)

Algorithm 11 Solve the optimization problem (4.5) using a greedy scheme. Inputs: the estimations of the coupled path cost for all the levels, $(\hat{\psi}_\ell)_{\ell=0}^{L}$, the estimation of the variance of the quantity of interest at level 0, $\hat{\mathcal{V}}_0$, the estimations of the differences of the quantity of interest for all the coupled levels, $(\hat{\mathcal{V}}_\ell)_{\ell=1}^{L}$, the prescribed tolerance, TOL, and the weak error estimation for level L, $\mathcal{E}_I$. Output: the number of realizations needed for each level, $(M)_{\ell=0}^{L}$.

Define $q_k := \dfrac{\sum_{\ell=0}^{L-k} \sqrt{\hat{\psi}_\ell \hat{\mathcal{V}}_\ell}}{RHS - \sum_{\ell=L-k+1}^{L} \hat{\mathcal{V}}_\ell}$

$RHS \leftarrow ((TOL - TOL^2 - \mathcal{E}_I)/C_A)^2$
fin $\leftarrow$ **false**
$k \leftarrow 0$
while not fin **and** $k \leq L$ **do**
  if $\hat{\psi}_{L-k} - q_k^2 \, \hat{\mathcal{V}}_{L-k} < 0$ **then**
    fin $\leftarrow$ **true**
    $(M_\ell)_{\ell=0}^{L-k} \leftarrow q_k \sqrt{\hat{\mathcal{V}}_\ell / \hat{\psi}_\ell}$
  else
    $M_{L-k} \leftarrow 1$
    $k \leftarrow k + 1$
  end if
end while
return $(M_\ell)_{\ell=0}^{L}$

Algorithm 12 Auxiliary function for Algorithm 7. Inputs: same as Algorithm 3. Outputs: the estimated runtime of the hybrid path at level 0, $\hat{\psi}_0$, an estimate of $\mathrm{Var}\left[g(\bar{X}(T))\right]$, $\mathcal{S}^2\left(g(\bar{X}(T);\cdot\right)$, an estimate of $\mathrm{E}\left[g(X(T))\right]$, $\mathcal{A}\left(g(\bar{X}(T);\cdot\right)$, an estimate of $\mathrm{E}\left[\mathcal{E}_I\right]$, $\mathcal{A}\left(\mathcal{E}_I;\cdot\right)$, an estimate of the expected number of steps needed by the SSA method, $\mathcal{A}\left(N_{\mathrm{SSA}^*};\cdot\right)$ and $\mathcal{A}\left(N_{\mathrm{TL}};\cdot\right)$. Here, $\mathbf{1}_{TL}(k) = 1$ if and only if the decision at time t_k was tau-leap. Notes: the values C_1, C_2 and C_3 are described in Section 4. Set appropriate values for M_0 and CV_0. For the sake of simplicity, we omit the arguments of the algorithms when there is no risk of confusion.

$M \leftarrow M_0$, $cv \leftarrow \infty$, $M_f \leftarrow 0$
while $cv > CV_0$ **do**
  for $m \leftarrow 1$ to M **do**
    $((\bar{X}(t_k))_{k=0}^K, N_{\mathrm{TL}}, N_{\mathrm{SSA},K1}, N_{\mathrm{SSA},K2}) \leftarrow$ generate one hybrid path (see [19])
    if the path does not exit $\mathbb{Z}_+^d$ **then**
      $M_f \leftarrow M_f + 1$
      Compute $g(\bar{X}(T;\bar{\omega}_m))$
      $\mathcal{E}_I \leftarrow$ Algorithm 9 with $(\bar{X}(t_k))_{k=0}^K$
      $(S_e(\bar{\omega}_m), S_v(\bar{\omega}_m)) \leftarrow$ Algorithm 10 with $(\bar{X}(t_k))_{k=0}^K$
      Estimate $N_{\mathrm{SSA}^*}(\bar{\omega}_m)$, using $\int_0^T a_0(\bar{\bar{X}}(s))ds$ (see [19])
      $C_{Poi}(\bar{\omega}_m) \leftarrow \sum_{j=1}^J \sum_{k=0}^K C_P(a_j(\bar{X}(t_k))(t_{k+1}-t_k))\mathbf{1}_{TL}(k)$
    end if
  end for
  $\hat{\mathcal{V}} \leftarrow \mathcal{S}^2\left(S_e; M_f\right) + \mathcal{A}\left(S_v; M_f\right)$
  Estimate the coefficients of variation $cv_{\mathcal{V}}$, cv_g and $cv_{\mathcal{E}_I}$ of the estimators of $\mathrm{Var}\left[g(\bar{X}(T))-g(\bar{\bar{X}}(T))\right]$, $\mathrm{Var}\left[g(\bar{X}(T))\right]$ and $\mathrm{E}\left[\mathcal{E}_I\right]$, respectively.
  $cv \leftarrow \max\{cv_{\mathcal{V}}, cv_g, cv_{\mathcal{E}_I}\}$
  $\hat{\psi}_0 \leftarrow C_1\mathcal{A}\left(N_{\mathrm{SSA},K1}; M_f\right) + C_2\mathcal{A}\left(N_{\mathrm{SSA},K2}; M_f\right) + C_3\mathcal{A}\left(N_{\mathrm{TL}}; M_f\right) + \mathcal{A}\left(C_{Poi}; M_f\right)$
  $M \leftarrow 2M$
end while
return $\left(\hat{\psi}_0, \mathcal{S}^2\left(g(\bar{X}(T)); M_f\right), \mathcal{A}\left(\{g(\bar{X}(T)), \mathcal{E}_I, N_{\mathrm{SSA}^*}, N_{\mathrm{TL}}\}; M_f\right)\right)$